

Product Manual

Seagate® Enterprise Performance 10K HDD v8 SAS

with TurboBoost™

| | Standard Models | Self-Encrypting Drive Models | SED FIPS140-2 Models Review Pending |
|---------------------------------------|-----------------|------------------------------|--|
| 512 Emulation | ST1800MM0018 | ST1800MM0068 | ST1800MM0078 |
| | ST1200MM0018 | ST1200MM0068 | |
| | ST900MM0018 | ST900MM0068 | |
| | ST600MM0018 | ST600MM0068 | |
| 512 Emulation w/TurboBoost | ST1800MM0128 | ST1800MM0148 | ST1800MM0158 |
| | ST1200MM0158 | ST1200MM0178 | |
| | ST900MM0128 | ST900MM0148 | |
| | ST600MM0158 | ST600MM0178 | |
| 4096 Native | ST1800MM0008 | ST1800MM0038 | ST1800MM0048 |
| | ST1200MM0008 | ST1200MM0038 | |
| | ST900MM0008 | ST900MM0038 | |
| | ST600MM0008 | ST600MM0038 | |
| 4096 Native w/TurboBoost | ST1800MM0088 | ST1800MM0108 | ST1800MM0118 |
| | ST1200MM0118 | ST1200MM0138 | |
| | ST900MM0088 | ST900MM0108 | |
| | ST600MM0118 | ST600MM0138 | |
| 512 Native | ST1200MM0088 | ST1200MM0098 | |
| | ST600MM0088 | ST600MM0098 | |

Document Revision History

| Revision | Date | Description of Change |
|----------|------------|-------------------------------------|
| Rev. A | 08/04/2014 | Initial release. |
| Rev. B | 10/27/2014 | fc, 7-9, 11-12, 14, 16, 18 & 54-59. |

© 2013 Seagate Technology LLC. All rights reserved.

Publication number: 100746003, Rev. B October 2014

Seagate, Seagate Technology and the Wave logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. Seagate Enterprise Performance 10K HDD and SeaTools are either trademarks or registered trademarks of Seagate Technology LLC or one of its affiliated companies in the United States and/or other countries. The FIPS logo is a certification mark of NIST, which does not imply product endorsement by NIST, the U.S., or Canadian governments. All other trademarks or registered trademarks are the property of their respective owners.

No part of this publication may be reproduced in any form without written permission of Seagate Technology LLC.

Call 877-PUB-TEK1 (877-782-8351) to request permission.

When referring to drive capacity, one gigabyte, or GB, equals one billion bytes and one terabyte, or TB, equals one trillion bytes. Your computer's operating system may use a different standard of measurement and report a lower capacity. In addition, some of the listed capacity is used for formatting and other functions, and thus will not be available for data storage. Actual quantities will vary based on various factors, including file size, file format, features and application software. Actual data rates may vary depending on operating environment and other factors. The export or re-export of hardware or software containing encryption may be regulated by the U.S. Department of Commerce, Bureau of Industry and Security (for more information, visit www.bis.doc.gov), and controlled for import and use outside of the U.S. Seagate reserves the right to change, without notice, product offerings or specifications.

| | |
|--|-----------|
| Seagate® Technology Support Services | 6 |
| 1.0 Scope | 7 |
| 2.0 Applicable standards and reference documentation | 8 |
| 2.1 Standards | 8 |
| 2.1.1 Electromagnetic compatibility | 8 |
| 2.1.2 Electromagnetic compliance | 9 |
| 2.1.3 European Union Restriction of Hazardous Substances (RoHS) | 9 |
| 2.1.4 China Restriction of Hazardous Substances (RoHS) Directive | 10 |
| 2.2 Reference documents | 10 |
| 3.0 General description | 11 |
| 3.1 Standard features | 11 |
| 3.2 Media description | 12 |
| 3.3 Performance | 12 |
| 3.4 Formatted capacities | 13 |
| 3.5 Programmable drive capacity | 13 |
| 3.6 Factory-installed options | 13 |
| 4.0 Performance characteristics | 14 |
| 4.1 Internal drive characteristics | 14 |
| 4.2 Performance characteristics | 14 |
| 4.2.1 Format command execution time (minutes) | 14 |
| 4.2.2 General performance characteristics | 14 |
| 4.3 Start/stop time | 15 |
| 4.4 Prefetch/multi-segmented cache control | 15 |
| 4.5 Cache operation | 15 |
| 4.5.1 Caching write data | 16 |
| 4.5.2 Prefetch operation | 16 |
| 4.5.3 TurboBoost caching operations | 16 |
| 5.0 Reliability specifications | 17 |
| 5.1 Error rates | 17 |
| 5.1.1 Recoverable Errors | 17 |
| 5.1.2 Unrecoverable Errors | 17 |
| 5.1.3 Seek errors | 17 |
| 5.1.4 Interface errors | 17 |
| 5.2 Reliability and service | 18 |
| 5.2.1 Annualized Failure Rate (AFR) and Mean Time Between Failure (MTBF) | 18 |
| 5.2.2 Preventive maintenance | 18 |
| 5.2.3 Hot plugging the drive | 18 |
| 5.2.4 S.M.A.R.T. | 18 |
| 5.2.5 Thermal monitor | 20 |
| 5.2.6 Drive Self Test (DST) | 20 |
| 5.2.7 Product warranty | 22 |
| 6.0 Physical/electrical specifications | 23 |
| 6.1 PowerChoice™ power management | 23 |
| 6.1.1 PowerChoice reporting methods | 24 |
| 6.2 AC power requirements | 24 |
| 6.3 DC power requirements | 24 |

CONTENTS

| | | |
|-------------|--|-----------|
| 6.3.1 | Conducted noise immunity | 29 |
| 6.3.2 | Power sequencing | 29 |
| 6.3.3 | Current profiles | 29 |
| 6.4 | Power dissipation | 33 |
| 6.5 | Environmental limits | 35 |
| 6.5.1 | Temperature | 35 |
| 6.5.2 | Humidity | 35 |
| 6.5.3 | Effective altitude (sea level) | 35 |
| 6.5.4 | Shock and vibration | 36 |
| 6.5.5 | Air cleanliness | 37 |
| 6.5.6 | Corrosive environment | 37 |
| 6.5.7 | Acoustics | 37 |
| 6.5.8 | Electromagnetic susceptibility | 37 |
| 6.6 | Mechanical specifications | 38 |
| <hr/> | | |
| 7.0 | About FIPS | 39 |
| <hr/> | | |
| 8.0 | About self-encrypting drives | 40 |
| 8.1 | Data encryption | 40 |
| 8.2 | Controlled access | 40 |
| 8.2.1 | Admin SP | 40 |
| 8.2.2 | Locking SP | 40 |
| 8.2.3 | Default password | 40 |
| 8.3 | Random number generator (RNG) | 41 |
| 8.4 | Drive locking | 41 |
| 8.5 | Data bands | 41 |
| 8.6 | Cryptographic erase | 41 |
| 8.7 | Authenticated firmware download | 42 |
| 8.8 | Power requirements | 42 |
| 8.9 | Supported commands | 42 |
| 8.10 | Sanitize - CRYPTOGRAPHIC ERASE | 42 |
| 8.11 | RevertSP | 42 |
| <hr/> | | |
| 9.0 | Defect and error management | 43 |
| 9.1 | Drive internal defects/errors | 43 |
| 9.2 | Drive error recovery procedures | 43 |
| 9.3 | SAS system errors | 44 |
| 9.4 | Background Media Scan | 45 |
| 9.5 | Deferred Auto-Reallocation | 45 |
| 9.6 | Idle Read After Write | 45 |
| 9.7 | Protection Information (PI) | 45 |
| 9.7.1 | Levels of PI | 46 |
| 9.7.2 | Setting and determining the current Type Level | 46 |
| 9.7.3 | Identifying a Protection Information drive | 46 |
| 9.8 | Seagate RAID Rebuild [™] | 46 |
| <hr/> | | |
| 10.0 | Installation | 47 |
| 10.1 | Drive orientation | 47 |
| 10.2 | Cooling | 47 |
| 10.3 | Drive mounting | 48 |
| 10.4 | Grounding | 48 |
| <hr/> | | |
| 11.0 | Interface requirements | 49 |
| 11.1 | SAS features | 49 |

CONTENTS

| | | |
|--------|---|----|
| 11.1.1 | Task management functions | 49 |
| 11.1.2 | Task management responses | 49 |
| 11.2 | Dual port support | 49 |
| 11.3 | SCSI commands supported | 50 |
| 11.3.1 | Inquiry data | 54 |
| 11.3.2 | Mode Sense data | 55 |
| 11.4 | Miscellaneous operating features and conditions | 60 |
| 11.4.1 | SAS physical interface | 60 |
| 11.4.2 | Physical characteristics | 63 |
| 11.4.3 | Connector requirements | 63 |
| 11.4.4 | Electrical description | 63 |
| 11.4.5 | Pin descriptions | 63 |
| 11.4.6 | SAS transmitters and receivers | 64 |
| 11.4.7 | Power | 64 |
| 11.5 | Signal characteristics | 64 |
| 11.5.1 | Ready LED Out | 64 |
| 11.5.2 | Differential signals | 65 |
| 11.6 | SAS-3 Specification compliance | 65 |
| 11.7 | Additional information | 65 |

| | | |
|------------|--|----|
| Figure 1. | Current profiles for 1800GB models | 29 |
| Figure 2. | Current profiles for 1200GB models. | 30 |
| Figure 3. | Current profiles for 900GB models. | 31 |
| Figure 4. | Current profiles for 600GB models. | 32 |
| Figure 5. | 1800GB (at 12Gb) DC current and power vs. input/output operations per second | 33 |
| Figure 6. | 1200GB (at 12Gb) DC current and power vs. input/output operations per second | 33 |
| Figure 7. | 900GB (at 12Gb) DC current and power vs. input/output operations per second | 34 |
| Figure 8. | 600GB (at 12Gb) DC current and power vs. input/output operations per second | 34 |
| Figure 9. | Recommended mounting. | 36 |
| Figure 10. | Mounting configuration dimensions | 38 |
| Figure 11. | Example of FIPS tamper evidence labels. | 39 |
| Figure 12. | Physical interface. | 47 |
| Figure 13. | Air flow | 48 |
| Figure 14. | Physical interface | 60 |
| Figure 15. | SAS device plug dimensions | 61 |
| Figure 16. | SAS device plug dimensions (detail) | 62 |
| Figure 17. | SAS transmitters and receivers | 64 |

Seagate® Technology Support Services

For information regarding online support and services, visit <http://www.seagate.com/about/contact-us/technical-support/>

Available services include:

- Presales & Technical support
- Global Support Services telephone numbers & business hours
- Authorized Service Centers

For information regarding Warranty Support, visit: <http://www.seagate.com/support/warranty-and-replacements/>

For information regarding data recovery services, visit <http://www.seagate.com/services-software/data-recovery-services/>

For Seagate OEM and Distribution partner portal, visit: <http://www.seagate.com/partners/>

For Seagate reseller portal, visit: <http://www.seagate.com/partners/my-spp-dashboard/>

1.0 Scope

This manual describes Seagate® Enterprise Performance 10K HDD v8 (Serial Attached SCSI) with TurboBoost™ disk drives.

Seagate Enterprise Performance 10K HDD drives support the SAS Protocol specifications to the extent described in this manual. The *SAS Interface Manual* (part number 100293071) describes the general SAS characteristics of this and other Seagate SAS drives. The Self-Encrypting Drive Reference Manual, part number 100515636, describes the interface, general operation, and security features available on Self-Encrypting Drive models.

Product data communicated in this manual is specific only to the model numbers listed in this manual. The data listed in this manual may not be predictive of future generation specifications or requirements. If integrators are designing a system which will use one of the models listed or future generation products and need further assistance, please contact the Field Applications Engineer (FAE) or our global support services group as shown on page 6.

Unless otherwise stated, the information in this manual applies to standard and Self-Encrypting Drive models.

| Standard Models | Self-Encrypting Drive (SED) Models | FIPS 140-2 Level 2 Models Review Pending |
|-----------------------------------|------------------------------------|--|
| 512 Emulation | | |
| ST1800MM0018 | ST1800MM0068 | ST1800MM0078 |
| ST1200MM0018 | ST1200MM0068 | |
| ST900MM0018 | ST900MM0068 | |
| ST600MM0018 | ST600MM0068 | |
| 512 Emulation w/TurboBoost | | |
| ST1800MM0128 | ST1800MM0148 | ST1200MM0158 |
| ST1200MM0158 | ST1200MM0178 | |
| ST900MM0128 | ST900MM0148 | |
| ST600MM0158 | ST600MM0178 | |
| 4096 Native | | |
| ST1800MM0008 | ST1800MM0038 | ST1800MM0048 |
| ST1200MM0008 | ST1200MM0038 | |
| ST900MM0008 | ST900MM0038 | |
| ST600MM0008 | ST600MM0038 | |
| 4096 Native w/TurboBoost | | |
| ST1800MM0088 | ST1800MM0108 | ST1800MM0118 |
| ST1200MM0118 | ST1200MM0138 | |
| ST900MM0088 | ST900MM0118 | |
| ST600MM0118 | ST600MM0138 | |
| 512 Native | | |
| ST1200MM0088 | ST1200MM0098 | |
| ST600MM0088 | ST600MM0098 | |

NOTE

Previous generations of Seagate Self-Encrypting Drive models were called Full Disk Encryption (FDE) models before a differentiation between drive-based encryption and other forms of encryption was necessary.

NOTE

The Self-Encrypting Drive models indicated on the cover of this product manual have provisions for "Security of Data at Rest" based on the standards defined by the Trusted Computing Group (see www.trustedcomputinggroup.org).

For more information on FIPS 140-2 Level 2 certification see Section 7.0 on page 39.

For product certification status visit - <http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401vend.htm>.

2.0 Applicable standards and reference documentation

The drives documented in this manual have been developed as system peripherals to the highest standards of design and construction. The drives depend on host equipment to provide adequate power and environment for optimum performance and compliance with applicable industry and governmental regulations. Special attention must be given in the areas of safety, power distribution, shielding, audible noise control, and temperature regulation. In particular, the drives must be securely mounted to guarantee the specified performance characteristics. Mounting by bottom holes must meet the requirements of Section 10.3.

2.1 Standards

The Seagate Enterprise Performance 10K HDD family complies with Seagate standards as noted in the appropriate sections of this manual and the Seagate *SAS Interface Manual*, part number 100293071.

The drives are recognized in accordance with UL 60950-1 as tested by UL, CSA 60950-1 as tested by CSA, and EN60950-1 as tested by TUV.

The security features of Self-Encrypting Drive models are based on the “TCG Storage Architecture Core Specification” and the “TCG Storage Workgroup Security Subsystem Class: Enterprise_A” specification with additional vendor-unique features as noted in this product manual.

2.1.1 Electromagnetic compatibility

The drive, as delivered, is designed for system integration and installation into a suitable enclosure prior to use. The drive is supplied as a subassembly and is not subject to Subpart B of Part 15 of the FCC Rules and Regulations.

The design characteristics of the drive serve to minimize radiation when installed in an enclosure that provides reasonable shielding. The drive is capable of meeting the Class B limits of the FCC Rules and Regulations when properly packaged; however, it is the user's responsibility to assure that the drive meets the appropriate EMI requirements in their system. Shielded I/O cables may be required if the enclosure does not provide adequate shielding. If the I/O cables are external to the enclosure, shielded cables should be used, with the shields grounded to the enclosure and to the host controller.

2.1.1.1 Electromagnetic susceptibility

As a component assembly, the drive is not required to meet any susceptibility performance requirements. It is the responsibility of those integrating the drive within their systems to perform those tests required and design their system to ensure that equipment operating in the same system as the drive or external to the system does not adversely affect the performance of the drive. See Section 6.3, DC power requirements.

2.1.2 Electromagnetic compliance

Seagate uses an independent laboratory to confirm compliance with the directives/standards for CE Marking and RCM Marking. The drive was tested in a representative system for typical applications and comply with the Electromagnetic Interference/Electromagnetic Susceptibility (EMI/EMS) for Class B products. The selected system represents the most popular characteristics for test platforms.

Although the test system with this Seagate model complies with the directives/standards, we cannot guarantee that all systems will comply. The computer manufacturer or system integrator shall confirm EMC compliance and provide the appropriate marking for their product.

Electromagnetic compliance for the European Union

If this model has the CE Marking it complies with the European Union requirements of the Electromagnetic Compatibility Directive 2004/108/EC as put into place on 20 July 2007.

Australian RCM

If this model has the RCM Marking it complies with the Australia/New Zealand Standard AS/NZ CISPR22 and meets the Electromagnetic Compatibility (EMC) Framework requirements of the Australian Communications and Media Authority (ACMA).

Canada ICES-003

If this model has the ICES-003 Marking it complies with the Canadian Standard Association Standard CAN/CSA-CISPR 22-10, Information Technology Equipment - Radio Disturbance Characteristics - Limits and Methods of Measurement.

Korean KCC

If these drives have the Korean Communications Commission (KCC) logo, they comply with KN22, KN 24, and KN61000.

Taiwanese BSMI

If this model has the Taiwanese certification mark then it complies with Chinese National Standard, CNS13438.

2.1.3 European Union Restriction of Hazardous Substances (RoHS)

The European Union Restriction of Hazardous Substances (RoHS) Directive restricts the presence of chemical substances, including Lead (Pb), in electronic products effective July 2006.

A number of parts and materials in Seagate products are procured from external suppliers. We rely on the representations of our suppliers regarding the presence of RoHS substances in these parts and materials. Our supplier contracts require compliance with our chemical substance restrictions, and our suppliers document their compliance with our requirements by providing material content declarations for all parts and materials for the disk drives documented in this publication. Current supplier declarations include disclosure of the inclusion of any RoHS-regulated substance in such parts or materials.

Seagate also has internal systems in place to ensure ongoing compliance with the RoHS Directive and all laws and regulations which restrict chemical content in electronic products. These systems include standard operating procedures that ensure that restricted substances are not utilized in our manufacturing operations, laboratory analytical validation testing, and an internal auditing process to ensure that all standard operating procedures are complied with.

2.1.4 China Restriction of Hazardous Substances (RoHS) Directive 中国限制危险物品的指令

This product has an Environmental Protection Use Period (EPUP) of 20 years. The following table contains information mandated by China's "Marking Requirements for Control of Pollution Caused by Electronic Information Products" Standard.



该产品具有20年的环境保护使用周期（EPUP）。下表包含了中国“电子产品所导致的污染的控制的记号要求”所指定的信息。

| Name of Parts 部件名称 | Toxic or Hazardous Substances or Elements有毒有害物质或元素 | | | | | |
|-----------------------|--|----------------------|----------------------|---|---|---|
| | Lead 铅 (Pb) | Mercury 汞 (Hg) | Cadmium 镉 (Cd) | Hexavalent Chromium 六价铬 (Cr6+) | Polybrominated Biphenyl 多溴联苯 (PBB) | Polybrominated Diphenyl Ether 多溴二苯醚 (PBDE) |
| PCBA | X | O | O | O | O | O |
| HDA | X | O | O | O | O | O |

"O" indicates the hazardous and toxic substance content of the part (at the homogeneous material level) is lower than the threshold defined by the China RoHS MCV Standard.

“O”表示该部件（于同类物品程度上）所含的危险和有毒物质低于中国RoHS MCV标准所定义的门槛值。

"X" indicates the hazardous and toxic substance content of the part (at the homogeneous material level) is over the threshold defined by the China RoHS MCV Standard.

“X”表示该部件（于同类物品程度上）所含的危险和有毒物质超出中国RoHS MCV标准所定义的门槛值。

2.2 Reference documents

| | |
|--|---|
| SCSI Commands Reference Manual | Seagate part number: 100293068 |
| SAS Interface Manual | Seagate part number: 100293071 |
| ANSI SAS Documents | |
| SFF-8223 | 2.5" Drive Form Factor with Serial Connector |
| SFF-8460 | HSS Backplane Design Guidelines |
| SFF-8470 | Multi Lane Copper Connector |
| SFF-8482 | SAS Plug Connector |
| ANSI INCITS.xxx | Serial Attached SCSI (SAS-3) Standard (T10//BSR INCITS 519 rev. 06) |
| ISO/IEC 14776-xxx | SCSI Architecture Model-3 (SAM-4) Standard (T10/1683-D) |
| ISO/IEC 14776-xxx | SCSI Primary Commands-3 (SPC-3) Standard (T10/1416-D) |
| ISO/IEC 14776-xxx | SCSI Block Commands-3 (SBC-2) Standard (T10/1417-D) |
| ANSI Small Computer System Interface (SCSI) Documents | |
| X3.270-1996 | (SCSI-3) Architecture Model |
| Trusted Computing Group (TCG) Documents (apply to Self-Encrypting Drive models only) | |
| TCG Storage Architecture Core Specification, Rev. 1.0 | |
| TCG Storage Security Subsystem Class Enterprise Specification, Rev. 1.0 | |
| Self-Encrypting Drives Reference Manual | Seagate part number: 100515636 |

In case of conflict between this document and any referenced document, this document takes precedence.

3.0 General description

Seagate® Enterprise Performance 10K HDD v8 drives provide high performance, high capacity data storage for a variety of systems including engineering workstations, network servers, mainframes, and supercomputers. The Serial Attached SCSI interface is designed to meet next-generation computing demands for performance, scalability, flexibility and high-density storage requirements.

The TurboBoost functionality provides improved performance over standard HDDs in real-world workloads. This improvement is due to the addition of a solid state component that caches “hot” data for reads as well as protects write data via non-volatile cache (NVC).

Seagate Enterprise Performance 10K HDD drives are random access storage devices designed to support the Serial Attached SCSI Protocol as described in the ANSI specifications, this document, and the *SAS Interface Manual* (part number 100293071) which describes the general interface characteristics of this drive. Seagate Enterprise Performance 10K HDD drives are classified as intelligent peripherals and provide level 2 conformance (highest level) with the ANSI SCSI-1 standard. The SAS connectors, cables and electrical interface are compatible with Serial ATA (SATA), giving future users the choice of populating their systems with either SAS or SATA hard disk drives. This allows integrators to continue to leverage existing investment in SCSI while gaining a 12Gb/s serial data transfer rate.

NOTE

Never disassemble the HDA and do not attempt to service items in the sealed enclosure (heads, media, actuator, etc.) as this requires special facilities. The drive does not contain user-replaceable parts. Opening the HDA for any reason voids the warranty.

Seagate Enterprise Performance 10K HDD drives use a dedicated load/unload zone at the outermost radius of the media to eliminate the possibility of destroying or degrading data by landing in the data zone. The heads automatically go to the ramp load/unload when power is removed from the drive and during the deeper sleep modes.

An automatic shipping lock prevents potential damage to the heads and discs that results from movement during shipping and handling. The shipping lock disengages and the head load process begins when power is applied to the drive.

The drives also use a high-performance actuator assembly with a low-inertia, balanced, patented, straight arm design that provides excellent performance with minimal power dissipation in any orientation.

3.1 Standard features

Seagate Enterprise Performance 10K HDD drives have the following standard features:

- 3.0 / 6.0 / 12.0 Gb Serial Attached SCSI (SAS) interface
- Integrated dual port SAS controller supporting the SCSI protocol
- Support for SAS expanders and fanout adapters
- Firmware downloadable using the SAS interface
- 128 - deep task set (queue)
- Supports up to 32 initiators
- 128MB data buffer (see Section 4.5)
- NVC-backed write cache of approximately 512K (4096 native and 512 emulation standard models)
- 32GB TurboBoost read cache
- 8MB TurboBoost NVC-backed write cache
- Jumperless configuration
- User-selectable logical block sizes for 512 emulation and native models (512, 520, 524 or 528 bytes per logical block)
- User-selectable logical block sizes for 4096 native models (4096, 4160, 4192 or 4224 bytes per logical block)
- Industry standard SFF 2.5-inch dimensions
- Programmable logical block reallocation scheme
- Flawed logical block reallocation at format time
- Programmable auto write and read reallocation
- Reallocation of defects on command (Post Format)
- ECC maximum burst correction length of 520 bits for 512 byte blocks and 1400 bits for 4K byte blocks.
- No preventive maintenance or adjustments required
- Dedicated head load/unload zone and automatic shipping lock
- Embedded servo design
- Self diagnostics performed when power is applied to the drive
- Zone bit recording (ZBR)
- Vertical, horizontal, or top down mounting

- Dynamic spindle brake
- Drive Self Test (DST)
- Background Media Scan (BMS)
- Idle Read After Write (IRAW)
- Power Choice
- SAS Power Disable
- RAID Rebuild TM

Seagate Enterprise Performance 10K HDD Self-Encrypting Drive models have the following additional features:

- Automatic data encryption/decryption
- Controlled access
- Random number generator
- Drive locking
- 16 independent data bands
- Cryptographic erase of user data for a drive that will be repurposed or scrapped
- Authenticated firmware download

3.2 Media description

The media used on the drive has an glass substrate coated with a thin film magnetic material, overcoated with a proprietary protective layer for improved durability and environmental protection.

3.3 Performance

- Firmware-controlled multisegmented cache designed to dynamically adjust segments for enhanced system performance
- 1200MB/s maximum instantaneous data transfers.
- 10K RPM spindle. Average latency = 2.9ms
- Background processing of queue
- Supports start and stop commands (spindle stops spinning)
- Adaptive seek velocity; improved seek performance

NOTE

There is no significant performance difference between Self-Encrypting Drive and standard (non-Self-Encrypting Drive) models.

3.4 Formatted capacities

Standard OEM models are formatted to 512 bytes per block for 512 emulation and native drives and 4096 bytes per block for 4096 native drives. The block size is selectable at format time. Supported block sizes are 512, 520, 524, and 528 for 512 emulation and native drives and 4096, 4160, 4192, and 4224 for 4096 native drives. Users having the necessary equipment may modify the data block size before issuing a format command and obtain different formatted capacities than those listed.

To provide a stable target capacity environment and at the same time provide users with flexibility if they choose, Seagate recommends product planning in one of two modes:

Seagate designs specify capacity points at certain block sizes that Seagate guarantees current and future products will meet. We recommend customers use this capacity in their project planning, as it ensures a stable operating point with backward and forward compatibility from generation to generation. The current guaranteed operating points for this product are shown below.

| Capacity (Blocks) | | | | | | | | |
|-------------------|---------------|-----------|---------------|-----------|---------------|-----------|---------------|-----------|
| Sector Size | 1800GB | | 1200GB | | 900GB | | 600GB | |
| | Decimal | Hex | Decimal | Hex | Decimal | Hex | Decimal | Hex |
| 512 | 3,516,328,368 | D196E9B0h | 2,344,225,968 | 8BBA0CB0h | 1,758,174,768 | 68CB9E30h | 1,172,123,568 | 45DD2FB0h |
| 520 | 3,441,923,088 | CD279410h | 2,294,615,392 | 88C50D60h | 1,720,961,544 | 6693CA08h | 1,147,307,696 | 446286B0h |
| 524 | 3,396,046,800 | CA6B8FD0h | 2,264,031,200 | 86F25FE0h | 1,698,023,400 | 6535C7E8h | 1,132,015,600 | 43792FF0h |
| 528 | 3,347,248,672 | C782F620h | 2,231,499,112 | 8501F968h | 1,673,624,336 | 63C17B10h | 1,115,749,560 | 4280FCB8h |
| 4096 | 439,541,046 | 1A32DD36h | 293,028,246 | 11774196h | 219,771,846 | D1973C6h | 146,515,446 | 8BBA5F6h |
| 4160 | 430,961,544 | 19AFF388h | 287,307,696 | 111FF7B0h | 215,480,776 | CD7F9C8h | 143,653,848 | 88FFBD8h |
| 4192 | 425,524,816 | 195CFE50h | 283,683,208 | 10E8A988h | 212,762,408 | CAE7F28h | 141,841,608 | 87454C8h |
| 4224 | 422,301,144 | 192BCDD8h | 281,534,096 | 10C7DE90h | 211,150,576 | C95E6F0h | 140,767,048 | 863EF48h |

3.5 Programmable drive capacity

Using the Mode Select command, the drive can change its capacity to something less than maximum. See the Mode Select (6) parameter list table in the *SAS Interface Manual*, part number 100293071. A value of zero in the Number of Blocks field indicates that the drive will not change the capacity it is currently formatted to have. A number other than zero and less than the maximum number of LBAs in the Number of Blocks field changes the total drive capacity to the value in the Number of Blocks field. A value greater than the maximum number of LBAs is rounded down to the maximum capacity.

3.6 Factory-installed options

The following items may be ordered which are incorporated at the manufacturing facility during production or packaged before shipping. Some of the options available are (not an exhaustive list of possible options):

- Other capacities can be ordered depending on sparing scheme and sector size requested.
- Single-unit shipping pack. The drive is normally shipped in bulk packaging to provide maximum protection against transit damage. Units shipped individually require additional protection as provided by the single unit shipping pack. Users planning single unit distribution should specify this option.
- The *Safety and Regulatory Agency Specifications*, part number 75789512, is usually included with each standard OEM drive shipped, but extra copies may be ordered.

4.0 Performance characteristics

This section provides detailed information concerning performance-related characteristics and features of Seagate Enterprise Performance 10K HDD drives.

4.1 Internal drive characteristics

| | 1800GB | 1200GB | 1200GB | 900GB | 600GB | 600GB | Models |
|----------------------------|----------------------|----------------------|------------|----------------------|----------------------|------------|---|
| Sector size | 4096 & 512 Emulation | 4096 & 512 Emulation | 512 Native | 4096 & 512 Emulation | 4096 & 512 Emulation | 512 Native | |
| Drive capacity | 1800 | 1200 | 1200 | 900 | 600 | 600 | GB (formatted, rounded off value) |
| Read/write data heads | 8 | 6 | 6 | 4 | 3 | 3 | |
| Bytes per track | 1065 | 1038 | 947 | 1065 | 1038 | 947 | Kbytes (avg, rounded off values) |
| Bytes per surface | 227,273 | 213,991 | 202,051 | 227,273 | 202,020 | 202,051 | MB (unformatted, rounded off value) |
| Tracks per surface (total) | 213.4 | 206.1 | 213.4 | 213.4 | 194.6 | 213.4 | Ktracks (user accessible, rounded off values) |
| Tracks per inch | 328 | 316.8 | 328 | 328 | 299 | 328 | KTPI (avg, rounded off values) |
| Peak bits per inch | 1993 | 1943 | 1952 | 1993 | 1943 | 1952 | Kb/in |
| Areal density | 644.6 | 606.9 | 619.85 | 644.6 | 573.0 | 619.9 | Gb/in ² |
| Disk rotation speed | 10K | 10K | 10K | 10K | 10K | 10K | rpm |
| Avg rotational latency | 2.9 | 2.9 | 2.9 | 2.9 | 2.9 | 2.9 | ms |

4.2 Performance characteristics

4.2.1 Format command execution time (minutes)

| 4096-byte sectors | 1800GB models | 1200GB models | 900GB models | 600GB models |
|--------------------------|---------------|---------------|--------------|--------------|
| Maximum (with verify) | 341 | 222 | 174 | 117 |
| Maximum (without verify) | 171 | 111 | 87 | 58 |

| 512-byte sectors | 1200GB models | 600GB models |
|--------------------------|---------------|--------------|
| Maximum (with verify) | 249 | 125 |
| Maximum (without verify) | 125 | 63 |

Execution time measured from receipt of the last byte of the Command Descriptor Block (CDB) to the request for a Status Byte Transfer to the Initiator (excluding connect/disconnect).

When changing sector sizes, the format times shown above may need to be increased by 30 minutes.

4.2.2 General performance characteristics

| | | |
|---|----------------|---|
| Sustained transfer rate for 4096 native and 512 emulation | All capacities | 112 to 230 MiB/s ** 117 to 241 MB/s |
| Sustained transfer rate for 512 native | 1.2TB & 600GB | 103 to 205 MiB/s ** 108 to 215 MB/s |
| SAS Interface maximum instantaneous transfer rate | | 1200 MB/s* per port (dual port = 2400 MB/s*) |
| Logical block sizes | | |
| 512 emulation and native - 512 (default), 520,524, or 528 | | |
| 4096 native - 4096 (default), 4160, 4192, or 4224 | | |
| Read/write consecutive sectors on a track | | Yes |
| Flaw reallocation performance impact (for flaws reallocated at format time using the spare sectors per sparing zone reallocation scheme.) | | Negligible |
| Average rotational latency | | 2.9ms |

*Assumes no errors and no relocated logical blocks. Rate measured from the start of the first logical block transfer to or from the host.

** MiB/s x 1.048 = MB/s

4.3 Start/stop time

The drive accepts the commands listed in the *SAS Interface Manual* less than 3 seconds after DC power has been applied.

If the drive receives a NOTIFY (ENABLE SPINUP) primitive through either port and has not received a START STOP UNIT command with the START bit equal to 0, the drive becomes ready for normal operations within 20 seconds (excluding the error recovery procedure).

If the drive receives a START STOP UNIT command with the START bit equal to 0 before receiving a NOTIFY (ENABLE SPINUP) primitive, the drive waits for a START STOP UNIT command with the START bit equal to 1. After receiving a START STOP UNIT command with the START bit equal to 1, the drive waits for a NOTIFY (ENABLE SPINUP) primitive. After receiving a NOTIFY (ENABLE SPINUP) primitive through either port, the drive becomes ready for normal operations within 20 seconds (excluding the error recovery procedure).

If the drive receives a START STOP UNIT command with the START bit and IMMED bit equal to 1 and does not receive a NOTIFY (ENABLE SPINUP) primitive within 5 seconds, the drive fails the START STOP UNIT command.

The START STOP UNIT command may be used to command the drive to stop the spindle. Stop time is 20 seconds (maximum) from removal of DC power. SCSI stop time is 20 seconds. There is no power control switch on the drive. However, power can be cycled on the drive by utilizing SAS Power Disable feature defined by T10 (i.e. drive Pin 3 high).

4.4 Prefetch/multi-segmented cache control

The drive provides a prefetch (read look-ahead) and multi-segmented cache control algorithms that in many cases can enhance system performance. Cache refers to the drive buffer storage space when it is used in cache operations. To select this feature, the host sends the Mode Select command with the proper values in the applicable bytes in page 08h. Prefetch and cache operations are independent features from the standpoint that each is enabled and disabled independently using the Mode Select command; however, in actual operation, the prefetch feature overlaps cache operation somewhat as described in sections 4.5.1 and 4.5.2.

All default cache and prefetch mode parameter values (Mode Page 08h) for standard OEM versions of this drive family are given in Section 11.3.2.

4.5 Cache operation

Note. Refer to the *SAS Interface Manual* for more detail concerning the cache bits.

Of the 128MB physical buffer space in the drive, approximately 60,000KB are available as a data cache. The remaining buffer space is reserved for internal drive use.

The drive keeps track of the logical block addresses of the data stored in each segment of the buffer. If the cache is enabled (see RCD bit in the *SAS Interface Manual*), data requested by the host with a read command is retrieved from the buffer, if possible, before any disk access is initiated. If cache operation is not enabled, the buffer is still used, but only as circular buffer segments during disk medium read operations (disregarding Prefetch operation for the moment). That is, the drive does not check in the buffer segments for the requested read data, but goes directly to the medium to retrieve it. The retrieved data merely passes through some buffer segment on the way to the host. All data transfers to the host are in accordance with buffer-full ratio rules. See the explanation provided with the information about Mode Page 02h (disconnect/reconnect control) in the *SAS Interface Manual*.

The following is a simplified description of the prefetch/cache operation:

Case A—read command is received and all of the requested logical blocks are already in the cache:

1. Drive transfers the requested logical blocks to the initiator.

Case B—A Read command requests data, and at least one requested logical block is not in any segment of the cache:

1. The drive fetches the requested logical blocks from the disk and transfers them into a segment, and then from there to the host in accordance with the Mode Select Disconnect/Reconnect parameters, page 02h.
2. If the prefetch feature is enabled, refer to section 4.5.2 for operation from this point.

Each cache segment is actually a self-contained circular buffer whose length is an integer number of logical blocks. The drive dynamically creates and removes segments based on the workload. The wrap-around capability of the individual segments greatly enhances the cache's overall performance.

NOTE

The size of each segment is not reported by Mode Sense command page 08h, bytes 14 and 15. The value 0XFFFF is always reported regardless of the actual size of the segment. Sending a size specification using the Mode Select command (bytes 14 and 15) does not set up a new segment size. If the STRICT bit in Mode page 00h (byte 2, bit 1) is set to one, the drive responds as it does for any attempt to change an unchangeable parameter.

4.5.1 Caching write data

NOTE

Write caching in this section is the traditional SCSI write caching (WCE=1) where writes are not protected on power loss. TurboBoost provides NVC-protected write caching when WCE=0 over a portion of the DRAM.

NOTE

Refer to the SAS Interface Manual for more detail concerning the cache bits.

Write caching is a write operation by the drive that makes use of a drive buffer storage area where the data to be written to the medium is stored while the drive performs the Write command.

If read caching is enabled (RCD=0), then data written to the medium is retained in the cache to be made available for future read cache hits. The same buffer space and segmentation is used as set up for read functions. The buffer segmentation scheme is set up or changed independently, having nothing to do with the state of RCD. When a write command is issued, if RCD=0, the cache is first checked to see if any logical blocks that are to be written are already stored in the cache from a previous read or write command. If there are, the respective cache segments are cleared. The new data is cached for subsequent Read commands.

If the number of write data logical blocks exceed the size of the segment being written into, when the end of the segment is reached, the data is written into the beginning of the same cache segment, overwriting the data that was written there at the beginning of the operation; however, the drive does not overwrite data that has not yet been written to the medium.

If write caching is enabled (WCE=1), then the drive may return Good status on a write command after the data has been transferred into the cache, but before the data has been written to the medium. If an error occurs while writing the data to the medium, and Good status has already been returned, a deferred error will be generated.

The Synchronize Cache command may be used to force the drive to write all cached write data to the medium. Upon completion of a Synchronize Cache command, all data received from previous write commands will have been written to the medium. Section 11.3.2 shows the mode default settings for the drive.

4.5.2 Prefetch operation

If the Prefetch feature is enabled, data in contiguous logical blocks on the disk immediately beyond that which was requested by a Read command are retrieved and stored in the buffer for immediate transfer from the buffer to the host on subsequent Read commands that request those logical blocks (this is true even if cache operation is disabled). Though the prefetch operation uses the buffer as a cache, finding the requested data in the buffer is a prefetch hit, not a cache operation hit.

To enable Prefetch, use Mode Select page 08h, byte 12, bit 5 (Disable Read Ahead - DRA bit). DRA bit = 0 enables prefetch.

The drive does not use the Max Prefetch field (bytes 8 and 9) or the Prefetch Ceiling field (bytes 10 and 11).

When prefetch (read look-ahead) is enabled (enabled by DRA = 0), the drive enables prefetch of contiguous blocks from the disk when it senses that a prefetch hit will likely occur. The drive disables prefetch when it decides that a prefetch hit is not likely to occur.

4.5.3 TurboBoost caching operations

Read data that has been promoted into the TurboBoost read cache does not persist through a power cycle. Read data is retrieved from the rotating media after a power cycle.

When WCE=0, TurboBoost provides NVC-protected write caching over the portion of the DRAM used to coalesce writes. Write data only goes into NVC when there is an unexpected power loss to the drive. The NVC has 90-day data retention. When WCE=1, a TurboBoost drive will operate on writes like a standard drive—writes in cache are not protected by NVC and may be lost with power loss.

5.0 Reliability specifications

The following reliability specifications assume correct host and drive operational interface, including all interface timings, power supply voltages, environmental requirements and drive mounting constraints.

| | |
|-----------------------------------|---|
| Seek error rate: | Less than 10 errors in 10^8 seeks |
| Read Error Rates ¹ | |
| Recovered Data | Less than 10 error in 10^{12} bits transferred (OEM default settings) |
| Unrecovered Data | Less than 1 sector in 10^{16} bits transferred |
| Miscorrected Data | Less than 1 sector in 10^{21} bits transferred |
| Interface error rate: | Less than 1 error in 10^{12} bits transferred |
| Mean Time Between Failure (MTBF): | 2,000,000 hours |
| Annualized Failure Rate (AFR): | 0.44% |
| Preventive maintenance: | None required |

1. Error rate specified with automatic retries and data correction with ECC enabled and all flaws reallocated.

5.1 Error rates

The error rates stated in this manual assume the following:

- The drive is operated in accordance with this manual using DC power as defined in paragraph 6.3, "DC power requirements."
- Errors caused by host system failures are excluded from error rate computations.
- Assume random data.
- Default OEM error recovery settings are applied. This includes AWRE, ARRE, full read retries, full write retries and full retry time.

5.1.1 Recoverable Errors

Recoverable errors are those detected and corrected by the drive, and do not require user intervention.

Recoverable Data errors will use Error Correction when needed.

Recovered Data error rate is determined using read bits transferred for recoverable errors occurring during a read, and using write bits transferred for recoverable errors occurring during a write.

5.1.2 Unrecoverable Errors

An unrecoverable data error is defined as a failure of the drive to recover data from the media. These errors occur due to head/media or write problems. Unrecoverable data errors are only detected during read operations, but not caused by the read. If an unrecoverable data error is detected, a MEDIUM ERROR (03h) in the Sense Key will be reported. Multiple unrecoverable data errors resulting from the same cause are treated as 1 error.

5.1.3 Seek errors

A seek error is defined as a failure of the drive to position the heads to the addressed track. After detecting an initial seek error, the drive automatically performs an error recovery process. If the error recovery process fails, a seek positioning error (Error code = 15h or 02h) will be reported with a Hardware error (04h) in the Sense Key. Recoverable seek errors are specified at Less than 10 errors in 10^8 seeks. Unrecoverable seek errors (Sense Key = 04h) are classified as drive failures.

5.1.4 Interface errors

An interface error is defined as a failure of the receiver on a port to recover the data as transmitted by the device port connected to the receiver. The error may be detected as a running disparity error, illegal code, loss of word sync, or CRC error.

5.2 Reliability and service

The reliability of Seagate Enterprise Performance 10K HDD disk drives can be enhanced by ensuring that the drive receives adequate cooling. Section 6.0 provides temperature measurements and other information that may be used to enhance the service life of the drive. Section 10.2 provides recommended air-flow information.

5.2.1 Annualized Failure Rate (AFR) and Mean Time Between Failure (MTBF)

The production disk drive shall achieve an annualized failure-rate of 0.44% (MTBF of 2,000,000 hours) over a 5 year service life when used in Enterprise Storage field conditions as limited by the following:

- 8760 power-on hours per year.
- HDA temperature as reported by the drive $\leq 50^{\circ}\text{C}$
- Ambient wet bulb temp $\leq 26^{\circ}\text{C}$
- Typical workload
- The AFR (MTBF) is a population statistic not relevant to individual units.
- ANSI/ISA S71.04-2013 G2 classification levels and dust contamination to ISO 14644-1 Class 8 standards (as measured at the device)

The MTBF specification for the drive assumes the operating environment is designed to maintain nominal drive temperature and humidity. Occasional excursions in operating conditions between the rated MTBF conditions and the maximum drive operating conditions may occur without significant impact to the rated MTBF. However continual or sustained operation beyond the rated MTBF conditions will degrade the drive MTBF and reduce product reliability.

5.2.2 Preventive maintenance

No routine scheduled preventive maintenance is required.

5.2.3 Hot plugging the drive

When a disk is powered on by switching the power or hot plugged, the drive runs a self test before attempting to communicate on its' interfaces. When the self test completes successfully, the drive initiates a Link Reset starting with OOB. An attached device should respond to the link reset. If the link reset attempt fails, or any time the drive loses sync, the drive initiates link reset. The drive will initiate link reset once per second but alternates between port A and B. Therefore each port will attempt a link reset once per 2 seconds assuming both ports are out of sync.

If the self-test fails, the drive does not respond to link reset on the failing port.

NOTE

It is the responsibility of the systems integrator to assure that no temperature, energy, voltage hazard, shorting of PCBA to ground, or ESD potential hazard is presented during the hot connect/disconnect operation. Discharge the static electricity from the drive carrier prior to inserting it into the system.

CAUTION

The drive motor must come to a complete stop prior to changing the plane of operation. This time is required to insure data integrity.

5.2.4 S.M.A.R.T.

S.M.A.R.T. is an acronym for Self-Monitoring Analysis and Reporting Technology. This technology is intended to recognize conditions that indicate imminent drive failure and is designed to provide sufficient warning of a failure to allow an application to back up the data before an actual failure occurs.

NOTE

The drive's firmware monitors specific attributes for degradation over time but can't predict instantaneous drive failures.

Each monitored attribute has been selected to monitor a specific set of failure conditions in the operating performance of the drive and the thresholds are optimized to minimize "false" and "failed" predictions.

Controlling S.M.A.R.T.

The operating mode of S.M.A.R.T. is controlled by the DEXCPT and PERF bits on the Informational Exceptions Control mode page (1Ch). Use the DEXCPT bit to enable or disable the S.M.A.R.T. feature. Setting the DEXCPT bit disables all S.M.A.R.T. functions. When enabled, S.M.A.R.T. collects on-line data as the drive performs normal read and write operations. When the PERF bit is set, the drive is considered to be in “On-line Mode Only” and will not perform off-line functions.

Applications can measure off-line attributes and force the drive to save the data by using the Rezero Unit command. Forcing S.M.A.R.T. resets the timer so that the next scheduled interrupt is in one hour.

Applications can interrogate the drive through the host to determine the time remaining before the next scheduled measurement and data logging process occurs. To accomplish this, issue a Log Sense command to log page 0x3E. This allows applications to control when S.M.A.R.T. interruptions occur. Forcing S.M.A.R.T. with the RTZ command resets the timer.

Performance impact

S.M.A.R.T. attribute data is saved to the disk so that the events that caused a predictive failure can be recreated. The drive measures and saves parameters once every hour subject to an idle period on the drive interfaces. The process of measuring off-line attribute data and saving data to the disk is interruptable. The maximum on-line only processing delay is summarized below

| | |
|---------------------------------|-------|
| Maximum processing delay | |
| Fully-enabled delay | |
| DEXCPT = 0 | |
| S.M.A.R.T. delay times | 75 ms |

Reporting control

Reporting is controlled by the MRIE bits in the Informational Exceptions Control mode page (1Ch). Subject to the reporting method. For example, if the MRIE is set to one, the firmware will issue to the host an 01-5D00 sense code. The FRU field contains the type of predictive failure that occurred. The error code is preserved through bus resets and power cycles.

Determining rate

S.M.A.R.T. monitors the rate at which errors occur and signals a predictive failure if the rate of degraded errors increases to an unacceptable level. To determine rate, error events are logged and compared to the number of total operations for a given attribute. The interval defines the number of operations over which to measure the rate. The counter that keeps track of the current number of operations is referred to as the Interval Counter.

S.M.A.R.T. measures error rates. All errors for each monitored attribute are recorded. A counter keeps track of the number of errors for the current interval. This counter is referred to as the Failure Counter.

Error rate is the number of errors per operation. The algorithm that S.M.A.R.T. uses to record rates of error is to set thresholds for the number of errors and their interval. If the number of errors exceeds the threshold before the interval expires, the error rate is considered to be unacceptable. If the number of errors does not exceed the threshold before the interval expires, the error rate is considered to be acceptable. In either case, the interval and failure counters are reset and the process starts over.

Predictive failures

S.M.A.R.T. signals predictive failures when the drive is performing unacceptably for a period of time. The firmware keeps a running count of the number of times the error rate for each attribute is unacceptable. To accomplish this, a counter is incremented each time the error rate is unacceptable and decremented (not to exceed zero) whenever the error rate is acceptable. If the counter continually increments such that it reaches the predictive threshold, a predictive failure is signaled. This counter is referred to as the Failure History Counter. There is a separate Failure History Counter for each attribute.

5.2.5 Thermal monitor

Seagate Enterprise Performance 10K HDD drives implement a temperature warning system which:

1. Signals the host if the temperature exceeds a value which would threaten the drive.
2. Saves a S.M.A.R.T. data frame on the drive which exceeds the threatening temperature value.

A temperature sensor monitors the drive temperature and issues a warning over the interface when the temperature exceeds a set threshold. The temperature is measured at power-up and then at ten-minute intervals after power-up.

The thermal monitor system generates a warning code of 01-0B01 when the temperature exceeds the specified limit in compliance with the SCSI standard.

This feature is controlled by the Enable Warning (EWasc) bit, and the reporting mechanism is controlled by the Method of Reporting Informational Exceptions field (MRIE) on the Informational Exceptions Control (IEC) mode page (1Ch).

5.2.6 Drive Self Test (DST)

Drive Self Test (DST) is a technology designed to recognize drive fault conditions that qualify the drive as a failed unit. DST validates the functionality of the drive at a system level.

There are two test coverage options implemented in DST:

1. Extended test
2. Short test

The most thorough option is the extended test that performs various tests on the drive and scans every logical block address (LBA) of the drive. The short test is time-restricted and limited in length—it does not scan the entire media surface, but does some fundamental tests and scans portions of the media.

If DST encounters an error during either of these tests, it reports a fault condition. If the drive fails the test, remove it from service and return it to Seagate for service.

5.2.6.1 DST failure definition

The drive will present a “diagnostic failed” condition through the self-tests results value of the diagnostic log page if a functional failure is encountered during DST. The channel and servo parameters are not modified to test the drive more stringently, and the number of retries are not reduced. All retries and recovery processes are enabled during the test. If data is recoverable, no failure condition will be reported regardless of the number of retries required to recover the data.

The following conditions are considered DST failure conditions:

- Seek error after retries are exhausted
- Track-follow error after retries are exhausted
- Read error after retries are exhausted
- Write error after retries are exhausted

Recovered errors will not be reported as diagnostic failures.

5.2.6.2 Implementation

This section provides all of the information necessary to implement the DST function on this drive.

5.2.6.2.1 State of the drive prior to testing

The drive must be in a ready state before issuing the Send Diagnostic command. There are multiple reasons why a drive may not be ready, some of which are valid conditions, and not errors. For example, a drive may be in process of doing a format, or another DST. It is the responsibility of the host application to determine the “not ready” cause.

While not technically part of DST, a Not Ready condition also qualifies the drive to be returned to Seagate as a failed drive.

A Drive Not Ready condition is reported by the drive under the following conditions:

- Motor will not spin
- Motor will not lock to speed
- Servo will not lock on track
- Drive cannot read configuration tables from the disk

In these conditions, the drive responds to a Test Unit Ready command with an 02/04/00 or 02/04/03 code.

5.2.6.2.2 Invoking DST

To invoke DST, submit the Send Diagnostic command with the appropriate Function Code (001b for the short test or 010b for the extended test) in bytes 1, bits 5, 6, and 7.

5.2.6.2.3 Short and extended tests

DST has two testing options:

1. short
2. extended

These testing options are described in the following two subsections.

Each test consists of three segments: an electrical test segment, a servo test segment, and a read/verify scan segment.

Short test (Function Code: 001b)

The purpose of the short test is to provide a time-limited test that tests as much of the drive as possible within 120 seconds. The short test does not scan the entire media surface, but does some fundamental tests and scans portions of the media. A complete read/verify scan is not performed and only factual failures will report a fault condition. This option provides a quick confidence test of the drive.

Extended test (Function Code: 010b)

The objective of the extended test option is to empirically test critical drive components. For example, the seek tests and on-track operations test the positioning mechanism. The read operation tests the read head element and the media surface. The write element is tested through read/write/read operations. The integrity of the media is checked through a read/verify scan of the media. Motor functionality is tested by default as a part of these tests.

The anticipated length of the Extended test is reported through the Control Mode page.

5.2.6.2.4 Log page entries

When the drive begins DST, it creates a new entry in the Self-test Results Log page. The new entry is created by inserting a new self-test parameter block at the beginning of the self-test results log parameter section of the log page. Existing data will be moved to make room for the new parameter block. The drive reports 20 parameter blocks in the log page. If there are more than 20 parameter blocks, the least recent parameter block will be deleted. The new parameter block will be initialized as follows:

1. The Function Code field is set to the same value as sent in the DST command
2. The Self-Test Results Value field is set to Fh
3. The drive will store the log page to non-volatile memory

After a self-test is complete or has been aborted, the drive updates the Self-Test Results Value field in its Self-Test Results Log page in non-volatile memory. The host may use Log Sense to read the results from up to the last 20 self-tests performed by the drive. The self-test results value is a 4-bit field that reports the results of the test. If the field is set to zero, the drive passed with no errors detected by the DST. If the field is not set to zero, the test failed for the reason reported in the field.

The drive will report the failure condition and LBA (if applicable) in the Self-test Results Log parameter. The Sense key, ASC, ASCQ, and FRU are used to report the failure condition.

5.2.6.2.5 Abort

There are multiple ways to abort a diagnostic. Two examples are: using a SCSI Bus Reset or a Bus Device message to abort the diagnostic

To abort a DST executing in background mode, use the abort code in the DST Function Code field. This will cause a 01 (self-test aborted by the application client) code to appear in the self-test results values log. All other abort mechanisms will be reported as a 02 (self-test routine was interrupted by a reset condition).

5.2.7 Product warranty

See “Seagate® Technology Support Services” on page 6. for warranty contact information.

Shipping

When transporting or shipping a drive, use only a Seagate-approved container. Keep the original box. Seagate approved containers are easily identified by the Seagate Approved Package label. Shipping a drive in a non-approved container voids the drive warranty.

Seagate repair centers may refuse receipt of components improperly packaged or obviously damaged in transit. Contact an authorized Seagate distributor to purchase additional boxes. Seagate recommends shipping by an air-ride carrier experienced in handling computer equipment.

Storage

The maximum recommended storage period for the drive in a non-operational environment is 90 days. Drives should be stored in the original unopened Seagate shipping packaging when ever possible. Once the drive is removed from the Seagate original packaging the recommended maximum period between drive operation cycles is 30 days. During any storage period the drive non-operational temperature, humidity, wet bulb, atmospheric conditions, shock, vibration, magnetic and electrical field specifications should be followed.

Product repair and return information

Seagate customer service centers are the only facilities authorized to service Seagate drives. Seagate does not sanction any third-party repair facilities. Any unauthorized repair or tampering with the factory seal voids the warranty.

6.0 Physical/electrical specifications

This section provides information relating to the physical and electrical characteristics of the drive.

6.1 PowerChoice™ power management

Drives using the load/unload architecture provide programmable power management to tailor systems for performance and greater energy efficiency.

The table below lists the supported PowerChoice modes. The further down in the table, the more power savings can be realized. For example, Idle_B mode results in greater power savings than Idle_A mode. Standby modes results in the greatest power savings.

PowerChoice modes

| Mode | Description |
|-----------|--|
| Idle_A | Reduced electronics |
| Idle_B | Heads unloaded. Disks spinning at full RPM |
| Idle_C | Heads unloaded. Disks spinning at reduced RPM |
| Standby_Y | Heads unloaded. Disks spinning at reduced RPM |
| Standby_Z | Heads unloaded. Motor stopped (disks not spinning) |

PowerChoice™ can be invoked using one of these two methods:

- Power Condition mode page method—Enable and initialize the idle condition timers and/or the standby condition timers. The timer values are based on the values set in the Power Condition mode page.
- START STOP UNIT command method—Use the START STOP UNIT command (OPERATION CODE 1Bh). This allows the host to directly transition the drive to any supported PowerChoice mode.

If both the Power Condition mode page and START STOP UNIT command methods are used, the START STOP UNIT command request takes precedence over the Power Condition mode page power control and may disable the idle condition and standby condition timers. The REQUEST SENSE command reports the current PowerChoice state if active and also the method by which the drive entered the PowerChoice state.

When the drive receives a command, all power condition timers are suspended if they were enabled via the Power Condition mode page. Once all outstanding commands are processed, the power condition timers are reinitialized to the values defined in the Power Condition mode page.

6.1.1 PowerChoice reporting methods

PowerChoice™ provides these reporting methods for tracking purposes:

Request Sense command reports

- Current power condition
- Method of entry

Note. Processing the Request Sense command does not impact the drive's power save state.

Mode Sense command reports (mode page 0x1A)

- Idle conditions enabled / disabled
- Idle condition timer values (100ms increments) (default, saved, current, changeable)

Power Condition Vital Product Data (VPD) Page (VPD page 0x8A)

- Supported power conditions
- Typical recovery time from power conditions (1ms increments)

Start/Stop Cycle Counter Log Page reports (log page 0x0E)

- Specified and accumulated Start/Stops and Load/Unload cycles

Power Condition Transitions Log Page reports (log page 0x1A, subpage 0x00)

- Accumulated transitions to Active, Idle_A, Idle_B, Idle_C, Standby_Y, Standby_Z

6.2 AC power requirements

None.

6.3 DC power requirements

The voltage and current requirements for a single drive are shown below. Values indicated apply at the drive connector.

The standard drive models and the SED drive models have identical hardware, however the security and encryption portion of the drive controller ASIC is enabled and functional in the SED models. This represents a small additional drain on the 5V supply of about 10mA and a commensurate increase of about 50mW in power consumption. There is no additional drain on the 12V supply.

Operation at 6.0Gb mode reduces the +5V supply load by 15mA with a commensurate power reduction of 75mW.

| | |
|-------------|--|
| NOTE | There is no measurable impact to the 12V supply load when running at lower interface speeds. |
|-------------|--|

Table 1 1800GB standard drive DC power requirements

| | | 12.0Gb mode | | |
|--------------------------------------|---------|-------------|----------|-------|
| | | (Amps) | (Amps) | Watts |
| | | | | Total |
| Voltage | | +5V | +12V [2] | |
| Regulation | [5] | ±5% | ±5% [2] | |
| Avg idle current DC | [1] [7] | 0.37 | 0.22 | 4.55 |
| Advanced Idle Current | | | | |
| Idle_A | | 0.40 | 0.24 | 4.91 |
| Idle_B | | 0.31 | 0.20 | 3.93 |
| Idle_C | | 0.30 | 0.07 | 2.31 |
| Standby | | 0.29 | 0.01 | 1.57 |
| Maximum starting current | | | | |
| (peak DC) DC | [3] | 0.61 | 0.89 | |
| (peak AC) AC | [3] | 0.72 | 1.42 | |
| Delayed motor start (max) DC | [1] [4] | 0.29 | 0.02 | 1.68 |
| Operating current (random read): | | | | |
| Typical DC | [1] [6] | 0.44 | 0.47 | 7.80 |
| Maximum DC | [1] | 0.49 | 0.54 | 8.90 |
| Maximum (peak) DC | | 0.77 | 1.08 | |
| Operating current (random write) | | | | |
| Typical DC | [1] | 0.44 | 0.42 | 7.22 |
| Maximum DC | [1] | 0.47 | 0.55 | 8.95 |
| Maximum (peak) DC | | 0.66 | 1.10 | |
| Operating current (sequential read) | | | | |
| Typical DC | [1] | 0.68 | 0.25 | 6.46 |
| Maximum DC | [1] | 0.75 | 0.28 | 7.05 |
| Maximum (peak) DC | | 1.30 | 0.38 | |
| Operating current (sequential write) | | | | |
| Typical DC | [1] | 0.62 | 0.28 | 6.44 |
| Maximum DC | [1] | 0.68 | 0.32 | 7.25 |
| Maximum (peak) DC | | 1.02 | 0.42 | |

Table 2 1200GB standard drive DC power requirements

| | | 12.0Gb mode | | |
|--------------------------------------|---------|-------------|----------|-------|
| | | (Amps) | (Amps) | Watts |
| | | | | Total |
| Voltage | | +5V | +12V [2] | |
| Regulation | [5] | ±5% | ±5% [2] | |
| Avg idle current DC | [1] [7] | 0.37 | 0.20 | 4.26 |
| Advanced Idle Current | | | | |
| Idle_A | | 0.40 | 0.23 | 4.80 |
| Idle_B | | 0.31 | 0.19 | 3.85 |
| Idle_C | | 0.30 | 0.07 | 2.31 |
| Standby | | 0.29 | 0.01 | 1.57 |
| Maximum starting current | | | | |
| (peak DC) DC | [3] | 0.61 | 1.26 | |
| (peak AC) AC | [3] | 0.72 | 1.40 | |
| Delayed motor start (max) DC | [1] [4] | 0.29 | 0.01 | 1.56 |
| Operating current (random read): | | | | |
| Typical DC | [1] [6] | 0.44 | 0.42 | 7.25 |
| Maximum DC | [1] | 0.49 | 0.53 | 8.85 |
| Maximum (peak) DC | | 0.77 | 1.07 | |
| Operating current (random write) | | | | |
| Typical DC | [1] | 0.44 | 0.38 | 6.78 |
| Maximum DC | [1] | 0.48 | 0.52 | 8.64 |
| Maximum (peak) DC | | 0.66 | 1.04 | |
| Operating current (sequential read) | | | | |
| Typical DC | [1] | 0.68 | 0.22 | 6.02 |
| Maximum DC | [1] | 0.75 | 0.28 | 7.09 |
| Maximum (peak) DC | | 1.30 | 0.33 | |
| Operating current (sequential write) | | | | |
| Typical DC | [1] | 0.62 | 0.24 | 5.97 |
| Maximum DC | [1] | 0.68 | 0.31 | 7.17 |
| Maximum (peak) DC | | 1.02 | 0.36 | |

Table 3 900GB standard drive DC power requirements

| | | 12.0Gb mode | | |
|--------------------------------------|---------|-------------|----------|-------|
| | | (Amps) | (Amps) | Watts |
| | | | | Total |
| Voltage | | +5V | +12V [2] | |
| Regulation | [5] | ±5% | ±5% [2] | |
| Avg idle current DC | [1] [7] | 0.37 | 0.17 | 3.88 |
| Advanced Idle Current | | | | |
| Idle_A | | 0.40 | 0.19 | 4.25 |
| Idle_B | | 0.31 | 0.16 | 3.43 |
| Idle_C | | 0.30 | 0.06 | 2.21 |
| Standby | | 0.29 | 0.01 | 1.58 |
| Maximum starting current | | | | |
| (peak DC) DC | [3] | 0.61 | 0.88 | |
| (peak AC) AC | [3] | 0.72 | 1.40 | |
| Delayed motor start (max) DC | [1] [4] | 0.29 | 0.01 | 1.56 |
| Operating current (random read): | | | | |
| Typical DC | [1] [6] | 0.44 | 0.41 | 7.12 |
| Maximum DC | [1] | 0.49 | 0.51 | 8.52 |
| Maximum (peak) DC | | 0.77 | 1.01 | |
| Operating current (random write) | | | | |
| Typical DC | [1] | 0.44 | 0.38 | 6.69 |
| Maximum DC | [1] | 0.48 | 0.47 | 8.10 |
| Maximum (peak) DC | | 0.66 | 0.95 | |
| Operating current (sequential read) | | | | |
| Typical DC | [1] | 0.68 | 0.20 | 5.81 |
| Maximum DC | [1] | 0.75 | 0.23 | 6.52 |
| Maximum (peak) DC | | 1.30 | 0.30 | |
| Operating current (sequential write) | | | | |
| Typical DC | [1] | 0.62 | 0.21 | 5.61 |
| Maximum DC | [1] | 0.68 | 0.28 | 6.77 |
| Maximum (peak) DC | | 1.02 | 0.32 | |

Table 4 600GB standard drive DC power requirements

| | | 12.0Gb mode | | |
|--------------------------------------|---------|-------------|----------|-------|
| | | (Amps) | (Amps) | Watts |
| | | | | Total |
| Voltage | | +5V | +12V [2] | |
| Regulation | [5] | ±5% | ±5% [2] | |
| Avg idle current DC | [1] [7] | 0.37 | 0.17 | 3.91 |
| Advanced Idle Current | | | | |
| Idle_A | | 0.40 | 0.19 | 4.31 |
| Idle_B | | 0.31 | 0.16 | 3.48 |
| Idle_C | | 0.30 | 0.06 | 2.22 |
| Standby | | 0.29 | 0.01 | 1.57 |
| Maximum starting current | | | | |
| (peak DC) DC | [3] | 0.61 | 0.86 | |
| (peak AC) AC | [3] | 0.72 | 1.44 | |
| Delayed motor start (max) DC | [1] [4] | 0.29 | 0.03 | 1.77 |
| Operating current (random read): | | | | |
| Typical DC | [1] [6] | 0.44 | 0.41 | 7.15 |
| Maximum DC | [1] | 0.49 | 0.50 | 8.44 |
| Maximum (peak) DC | | 0.77 | 1.00 | |
| Operating current (random write) | | | | |
| Typical DC | [1] | 0.44 | 0.39 | 6.84 |
| Maximum DC | [1] | 0.48 | 0.47 | 8.05 |
| Maximum (peak) DC | | 0.66 | 0.94 | |
| Operating current (sequential read) | | | | |
| Typical DC | [1] | 0.68 | 0.19 | 5.69 |
| Maximum DC | [1] | 0.75 | 0.24 | 6.60 |
| Maximum (peak) DC | | 1.30 | 0.28 | |
| Operating current (sequential write) | | | | |
| Typical DC | [1] | 0.62 | 0.21 | 5.64 |
| Maximum DC | [1] | 0.68 | 0.27 | 6.62 |
| Maximum (peak) DC | | 1.02 | 0.32 | |

- [1] Measured with average reading DC ammeter. Instantaneous +12V current peaks will exceed these values. Power supply at nominal voltage. N (number of drives tested) = 6, 35 Degrees C ambient.
- [2] For +12 V, a -10% tolerance is allowed during initial spindle start but must return to ±5% before reaching 9,936 RPM. The ±5% must be maintained after the drive signifies that its power-up sequence has been completed and that the drive is able to accept selection by the host initiator.
- [3] See +12V current profile in Figure 1 (for 1200GB models).
- [4] This condition occurs after OOB and Speed Negotiation completes but before the drive has received the Notify Spinup primitive.
- [5] See paragraph 6.3.1, "Conducted noise immunity." Specified voltage tolerance includes, noise, and transient response.
- [6] Operating condition is defined as random 8 block reads.
- [7] During idle, the drive heads are relocated every 30 seconds to a random location within the band from three-quarters to maximum track.

General DC power requirement notes.

1. Minimum current loading for each supply voltage is not less than 1.7% of the maximum operating current shown.
2. The +5V and +12V supplies should employ separate ground returns.
3. Where power is provided to multiple drives from a common supply, careful consideration for individual drive power requirements should be noted. Where multiple units are powered on simultaneously, the peak starting current must be available to each device.
4. Parameters, other than spindle start, are measured after a 10-minute warm up.
5. No terminator power.

6.3.1 Conducted noise immunity

Noise is specified as a periodic and random distribution of frequencies covering a defined frequency range. Maximum allowed noise values given below are peak-to-peak measurements and apply at the drive power connector.

| | | |
|------|---|-----------------------------------|
| +5v | = | 250 mV pp from 100 Hz to 20 MHz. |
| +12v | = | 450 mV pp from 100 Hz to 100 KHz. |
| | | 250 mV pp from 100 KHz to 20 MHz. |
| | | 150 mV pp from 20 MHz to 80 MHz. |

6.3.2 Power sequencing

The drive does not require power sequencing. The drive protects against inadvertent writing during power-up and down.

6.3.3 Current profiles

The +12V and +5V current profiles for the Seagate Enterprise Performance 10K HDD drives are shown below.




Figure 1. Current profiles for 1800GB models




Figure 2. Current profiles for 1200GB models




Figure 3. Current profiles for 900GB models




Figure 4. Current profiles for 600GB models

Note: All times and currents are typical. See Section 6.3 for maximum current requirements.

6.4 Power dissipation

1800GB models in 12Gb operation

Typical power dissipation under idle conditions in 12Gb operation is 4.55 watts 15.52 BTUs per hour).

To obtain operating power for typical random read operations, refer to the following I/O rate curve (see Figure 6). Locate the typical I/O rate for a drive in the system on the horizontal axis and read the corresponding +5 volt current, +12 volt current, and total watts on the vertical axis. To calculate BTUs per hour, multiply watts by 3.4123.




Figure 5. 1800GB (at 12Gb) DC current and power vs. input/output operations per second

1200GB models in 12Gb operation

Typical power dissipation under idle conditions in 12Gb operation is 4.26 watts 14.54 BTUs per hour).

To obtain operating power for typical random read operations, refer to the following I/O rate curve (see Figure 6). Locate the typical I/O rate for a drive in the system on the horizontal axis and read the corresponding +5 volt current, +12 volt current, and total watts on the vertical axis. To calculate BTUs per hour, multiply watts by 3.4123.




Figure 6. 1200GB (at 12Gb) DC current and power vs. input/output operations per second

900GB models in 12Gb operation

Typical power dissipation under idle conditions in 12Gb operation is 3.88 watts 13.25 BTUs per hour).

To obtain operating power for typical random read operations, refer to the following I/O rate curve (see Figure 6). Locate the typical I/O rate for a drive in the system on the horizontal axis and read the corresponding +5 volt current, +12 volt current, and total watts on the vertical axis. To calculate BTUs per hour, multiply watts by 3.4123.




Figure 7. 900GB (at 12Gb) DC current and power vs. input/output operations per second

600GB models in 12Gb operation

Typical power dissipation under idle conditions in 12Gb operation is 3.91 watts 13.35 BTUs per hour).

To obtain operating power for typical random read operations, refer to the following I/O rate curve (see Figure 6). Locate the typical I/O rate for a drive in the system on the horizontal axis and read the corresponding +5 volt current, +12 volt current, and total watts on the vertical axis. To calculate BTUs per hour, multiply watts by 3.4123.




Figure 8. 600GB (at 12Gb) DC current and power vs. input/output operations per second

6.5 Environmental limits

Temperature and humidity values experienced by the drive must be such that condensation does not occur on any drive part. Altitude and atmospheric pressure specifications are referenced to a standard day at 58.7°F (14.8°C).

| | |
|-------------|--|
| NOTE | To maintain optimal performance drives should be run at nominal drive temperatures and humidity. |
|-------------|--|

6.5.1 Temperature

a. Operating

41°F to 131°F (5°C to 55°C) drive temperature range with a maximum temperature gradient of 36°F (20°C) per hour.

The maximum allowable drive temperature is 140°F (60°C).

| | |
|-------------|--|
| NOTE | Operating drive temperature references temperature reported by the drive on Log Page 0Dh |
|-------------|--|

| | |
|-------------|---|
| NOTE | Air flow may be required to achieve consistent nominal drive temperature values (see Section 9.2). To confirm that the required cooling is provided for the electronics and HDA, place the drive in its final mechanical configuration, and perform random write/read operations. After the temperatures stabilize, read drive temperature from Temperature Log Page (0Dh). |
|-------------|---|

b. Non-operating

–40° to 158°F (–40° to 70°C) package ambient with a maximum gradient of 36°F (20°C) per hour. This specification assumes that the drive is packaged in the shipping container designed by Seagate for use with drive.

6.5.2 Humidity

The values below assume that no condensation on the drive occurs. Maximum wet bulb temperature is 84.2°F (29°C).

a. Operating

5% to 95% non-condensing relative humidity with a maximum gradient of 20% per hour.

b. Non-operating

5% to 95% non-condensing relative humidity with a maximum gradient of 20% per hour.

6.5.3 Effective altitude (sea level)

a. Operating

–1000 to +10,000 feet (–304.8 to +3048 meters)

b. Non-operating

–1000 to +40,000 feet (–304.8 to +12,192 meters)

6.5.4 Shock and vibration

Shock and vibration limits specified in this document are measured directly on the drive chassis. If the drive is installed in an enclosure to which the stated shock and/or vibration criteria is applied, resonances may occur internally to the enclosure resulting in drive movement in excess of the stated limits. If this situation is apparent, it may be necessary to modify the enclosure to minimize drive movement.

The limits of shock and vibration defined within this document are specified with the drive mounted by any of the four methods shown in Figure 9, and in accordance with the restrictions of Section 10.3.

6.5.4.1 Shock

a. Operating

The drive, as installed for normal operation, shall operate error free while subjected to intermittent shock not exceeding:

- 40 Gs at a maximum duration of 11ms (half sinewave)
- 25 Gs at a maximum duration of 2ms (half sinewave)

Shock may be applied in the X, Y, or Z axis. Shock is not to be repeated more than once every 2 seconds.

b. Non-operating

The limits of non-operating shock shall apply to all conditions of handling and transportation. This includes both isolated drives and integrated drives.

The drive subjected to nonrepetitive shock not exceeding the three values below, shall not exhibit device damage or performance degradation.

- 80 Gs at a maximum duration of 11ms (half sinewave)
- 400 Gs at a maximum duration of 2ms (half sinewave)
- 300 Gs at a maximum duration of 0.5ms (half sinewave)

Shock may be applied in the X, Y, or Z axis.

c. Packaged

Seagate finished drive bulk packs are designed and tested to meet or exceed applicable ISTA and ASTM standards. Volume finished drives will be shipped from Seagate factories on pallets to minimize freight costs and ease material handling. Seagate finished drive bulk packs may be shipped individually. For less than full shipments, instructions are printed on the bulk pack carton for minimum drive quantities and proper drive placement.




Figure 9. Recommended mounting

NOTE Image is for reference only, may not represent actual drive.

6.5.4.2 Vibration

a. Operating—normal

The drive as installed for normal operation, shall comply with the complete specified performance while subjected to continuous vibration not exceeding

5-500 Hz @ 0.5 G (zero to peak)

Vibration may be applied in the X, Y, or Z axis.

Operating normal translational random flat profile

10 - 500 Hz (translational random flat profile) 0.5 GRMS

b. Non-operating

The limits of non-operating vibration shall apply to all conditions of handling and transportation. This includes both isolated drives and integrated drives.

The drive shall not incur physical damage or degraded performance as a result of continuous vibration not exceeding 5-500 Hz @ 3.0 G (zero to peak)

Vibration may be applied in the X, Y, or Z axis.

Non-operating translational random flat profile

10 - 500 Hz (translational random flat profile) 2.4 GRMS

6.5.5 Air cleanliness

The drive is designed to operate in a typical office environment with minimal environmental control.

6.5.6 Corrosive environment

Seagate electronic drive components pass accelerated corrosion testing equivalent to 10 years exposure to light industrial environments containing sulfurous gases, chlorine and nitric oxide, classes G and H per ASTM B845. However, this accelerated testing cannot duplicate every potential application environment.

Users should use caution exposing any electronic components to uncontrolled chemical pollutants and corrosive chemicals as electronic drive component reliability can be affected by the installation environment. The silver, copper, nickel and gold films used in hard disk drives are especially sensitive to the presence of sulfide, chloride, and nitrate contaminants. Sulfur is found to be the most damaging. Materials used in cabinet fabrication, such as vulcanized rubber, that can outgas corrosive compounds should be minimized or eliminated. The useful life of any electronic equipment may be extended by replacing materials near circuitry with sulfide-free alternatives.

Seagate recommends that data centers be kept clean by monitoring and controlling the dust and gaseous contamination. Gaseous contamination should be within ANSI/ISA S71.04-2013 G2 classification levels (as measured on copper and silver coupons), and dust contamination to ISO 14644-1 Class 8 standards, and MTBF rated conditions as defined in the Annualized Failure Rate (AFR) and Mean Time Between Failure (MTBF) section.

6.5.7 Acoustics

Sound power during idle mode shall be 3.1 bels typical when measured to ISO 7779 specification.

There will not be any discrete tones more than 9 dB above the masking noise on typical drives when measured according to Seagate specification 30553-001. There will not be any tones more than 24 dB above the masking noise on any drive.

6.5.8 Electromagnetic susceptibility

See Section 2.1.1.1.

6.6 Mechanical specifications

Refer to Figure 10 for detailed mounting configuration dimensions. See Section 10.3, “Drive mounting.”

Weight:

| | | |
|---------------|-------------|----------------|
| 1800GB models | .439 pounds | .199 kilograms |
| 1200GB models | .426 pounds | .193 kilograms |
| 900GB models | | |
| 600GB models | .417 pounds | .189 kilograms |

NOTE

These dimensions conform to the Small Form Factor Standard documented in SFF-8201 and SFF-8223, found at www.sffcommittee.org




Figure 10. Mounting configuration dimensions

7.0 About FIPS

The Federal Information Processing Standard (FIPS) Publication 140-2 is a U.S. Government Computer Security Standard used to accredit cryptographic modules. It is titled 'Security Requirements for Cryptographic Modules (FIPS PUB 140-2)' and is issued by the National Institute of Standards and Technology (NIST).

Purpose

This standard specifies the security requirements that will be satisfied by a cryptographic module utilized within a security system protecting sensitive but unclassified information. The standard provides four increasing, qualitative levels of security: Level 1, Level 2, Level 3 and Level 4. These levels are intended to cover the wide range of potential applications and environments in which cryptographic modules may be employed.

Validation Program

Products that claim conformance to this standard are validated by the Cryptographic Module Validation Program (CMVP) which is a joint effort between National Institute of Standards and Technology (NIST) and the Communications Security Establishment (CSE) of the Government of Canada. Products validated as conforming to FIPS 140-2 are accepted by the Federal agencies of both countries for the protection of sensitive information (United States) or Designated Information (Canada).

In the CMVP, vendors of cryptographic modules use independent, accredited testing laboratories to have their modules tested. National Voluntary Laboratory Accreditation Program (NVLAP) accredited laboratories perform cryptographic module compliance/conformance testing.

Seagate Enterprise SED

The SEDs referenced in this Product Manual have been validated by CMVP and have been thoroughly tested by a NVLAP accredited lab to satisfy FIPS 140-2 Level 2 requirements. In order to operate in FIPS Approved Mode of Operation, these SEDs require security initialization. For more information, refer to 'Security Rules' section in the 'Security Policy' document uploaded on the NIST website. To reference the product certification visit - <http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401vend.htm>, and search for "Seagate".

Security Level 2

Security Level 2 enhances the physical security mechanisms of a Security Level 1 cryptographic module by adding the requirement for tamper-evidence, which includes the use of tamper-evident coatings or seals on removable covers of the module. Tamper-evident coatings or seals are placed on a cryptographic module so that the coating or seal must be broken to attain physical access to the critical security parameters (CSP) within the module. Tamper-evident seals (example shown in Figure 11, page 39) are placed on covers to protect against unauthorized physical access. In addition Security Level 2 requires, at a minimum, role-based authentication in which a cryptographic module authenticates the authorization of an operator to assume a specific role and perform a corresponding set of services.



Figure 11. Example of FIPS tamper evidence labels.

| | |
|-------------|--|
| NOTE | Image is for reference only, may not represent actual drive. |
|-------------|--|

8.0 About self-encrypting drives

Self-encrypting drives (SEDs) offer encryption and security services for the protection of stored data, commonly known as “protection of data at rest.” These drives are compliant with the Trusted Computing Group (TCG) Enterprise Storage Specifications as detailed in Section 2.2.

The Trusted Computing Group (TCG) is an organization sponsored and operated by companies in the computer, storage and digital communications industry. Seagate’s SED models comply with the standards published by the TCG.

To use the security features in the drive, the host must be capable of constructing and issuing the following two SCSI commands:

- Security Protocol Out
- Security Protocol In

These commands are used to convey the TCG protocol to and from the drive in their command payloads.

8.1 Data encryption

Encrypting drives use one inline encryption engine for each port, employing AES 256-bit data encryption keys with AES-XTS mode to encrypt all data prior to being written on the media and to decrypt all data as it is read from the media. The encryption engines are always in operation and cannot be disabled.

The 32-byte Data Encryption Key (DEK) is a random number which is generated by the drive, never leaves the drive, and is inaccessible to the host system. The DEK is itself encrypted when it is stored on the media and when it is in volatile temporary storage (DRAM) external to the encryption engine. A unique data encryption key is used for each of the drive's possible 16 data bands (see Section 8.5).

8.2 Controlled access

The drive has two security providers (SPs) called the "Admin SP" and the "Locking SP". These act as gatekeepers to the drive security services. Security-related commands will not be accepted unless they also supply the correct credentials to prove the requester is authorized to perform the command.

8.2.1 Admin SP

The Admin SP allows the drive's owner to enable or disable firmware download operations (see Section 8.4). Access to the Admin SP is available using the SID (Secure ID) password or the MSID (Manufacturers Secure ID) password.

8.2.2 Locking SP

The Locking SP controls read/write access to the media and the cryptographic erase feature. Access to the Locking SP is available using the BandMasterX or EraseMaster passwords. Since the drive owner can define up to 16 data bands on the drive, each data band has its own password called BandMasterX where X is the number of the data band (0 through 15).

8.2.3 Default password

When the drive is shipped from the factory, all passwords are set to the value of MSID. This 32-byte random value can only be read by the host electronically over the interface. After receipt of the drive, it is the responsibility of the owner to use the default MSID password as the authority to change all other passwords to unique owner-specified values.

8.3 Random number generator (RNG)

The drive has a 32-byte hardware RNG that it uses to derive encryption keys or, if requested to do so, to provide random numbers to the host for system use, including using these numbers as Authentication Keys (passwords) for the drive's Admin and Locking SPs.

8.4 Drive locking

In addition to changing the passwords, as described in Section 8.2.3, the owner should also set the data access controls for the individual bands.

The variable "LockOnReset" should be set to "PowerCycle" to ensure that the data bands will be locked if power is lost. In addition "ReadLockEnabled" and "WriteLockEnabled" must be set to true in the locking table in order for the bands "LockOnReset" setting of "PowerCycle" to actually lock access to the band when a "PowerCycle" event occurs. This scenario occurs if the drive is removed from its cabinet. The drive will not honor any data read or write requests until the bands have been unlocked. This prevents the user data from being accessed without the appropriate credentials when the drive has been removed from its cabinet and installed in another system.

When the drive is shipped from the factory, the firmware download port is unlocked.

8.5 Data bands

When shipped from the factory, the drive is configured with a single data band called Band 0 (also known as the Global Data Band) which comprises LBA 0 through LBA max. The host may allocate Band1 by specifying a start LBA and an LBA range. The real estate for this band is taken from the Global Band. An additional 14 Data Bands may be defined in a similar way (Band2 through Band15) but before these bands can be allocated LBA space, they must first be individually enabled using the EraseMaster password.

Data bands cannot overlap but they can be sequential with one band ending at LBA (x) and the next beginning at LBA (x+1).

Each data band has its own drive-generated encryption key and its own user-supplied password. The host may change the Encryption Key (see Section 8.6) or the password when required. The bands should be aligned to 4096 native LBA boundaries.

8.6 Cryptographic erase

A significant feature of SEDs is the ability to perform a cryptographic erase. This involves the host telling the drive to change the data encryption key for a particular band. Once changed, the data is no longer recoverable since it was written with one key and will be read using a different key. Since the drive overwrites the old key with the new one, and keeps no history of key changes, the user data can never be recovered. This is tantamount to an instantaneous data erase and is very useful if the drive is to be scrapped or redispositioned.

8.7 Authenticated firmware download

In addition to providing a locking mechanism to prevent unwanted firmware download attempts, the drive also only accepts download files which have been cryptographically signed by the appropriate Seagate Design Center.

Three conditions must be met before the drive will allow the download operation:

1. The download must be an SED file. A standard (base) drive (non-SED) file will be rejected.
2. The download file must be signed and authenticated.
3. As with a non-SED drive, the download file must pass the acceptance criteria for the drive. For example it must be applicable to the correct drive model, and have compatible revision and customer status.

8.8 Power requirements

The standard drive models and the SED drive models have identical hardware, however the security and encryption portion of the drive controller ASIC is enabled and functional in the SED models. This represents a small additional drain on the 5V supply of about 30mA and a commensurate increase of about 150mW in power consumption. There is no additional drain on the 12V supply. See the tables in Section 6.3 for power requirements on the standard (non-SED) drive models.

8.9 Supported commands

The SED models support the following two commands in addition to the commands supported by the standard (non-SED) models as listed in Section 8.9, "Supported commands":

- Security Protocol Out (B5h)
- Security Protocol In (A2h)

8.10 Sanitize - CRYPTOGRAPHIC ERASE

This command cryptographically erases all user data on the drive by destroying the current data encryption key and replacing it with a new data encryption key randomly generated by the drive. Sanitize CRYPTOGRAPHIC ERASE is a SCSI CDB Op code 48h and selecting the service action code 3 (CRYPTOGRAPHIC ERASE)

8.11 RevertSP

The SED models will support the RevertSP feature which erases all data in all bands on the device and returns the contents of all SPs (Security Providers) on the device to their original factory state. In order to execute the RevertSP method the unique PSID (Physical Secure ID) printed on the drive label must be provided. PSID is not electronically accessible and can only be manually read from the drive label or scanned in via the 2D barcode.

9.0 Defect and error management

Seagate continues to use innovative technologies to manage defects and errors. These technologies are designed to increase data integrity, perform drive self-maintenance, and validate proper drive operation.

SCSI defect and error management involves drive internal defect/error management and SAS system error considerations (errors in communications between the initiator and the drive). In addition, Seagate provides the following technologies used to increase data integrity and drive reliability:

- Background Media Scan (see Section 9.4)
- Deferred Auto-Reallocation (see Section 9.5)
- Idle Read After Write (see Section 9.6)

The read error rates and specified storage capacities are not dependent on host (initiator) defect management routines.

9.1 Drive internal defects/errors

During the initial drive format operation at the factory, media defects are identified, tagged as being unusable, and their locations recorded on the drive primary defects list (referred to as the “P” list and also as the ETF defect list). At factory format time, these known defects are also reallocated, that is, reassigned to a new place on the medium and the location listed in the defects reallocation table. The “P” list is not altered after factory formatting. Locations of defects found and reallocated during error recovery procedures after drive shipment are listed in the “G” list (defects growth list). The “P” and “G” lists may be referenced by the initiator using the Read Defect Data command.

Details of the SCSI commands supported by the drive are described in the *SAS Interface Manual*. Also, more information on the drive Error Recovery philosophy is presented in the *SAS Interface Manual*.

9.2 Drive error recovery procedures

When an error occurs during drive operation, the drive, if programmed to do so, performs error recovery procedures to attempt to recover the data. The error recovery procedures used depend on the options previously set in the Error Recovery Parameters mode page. Error recovery and defect management may involve using several SCSI commands described in the *SCSI Interface Manual*. The drive implements selectable error recovery time limits required in video applications.

The error recovery scheme supported by the drive provides a way to control the total error recovery time for the entire command in addition to controlling the recovery level for a single LBA. The total amount of time spent in error recovery for a command can be limited using the Recovery Time Limit bytes in the Error Recovery mode page. The total amount of time spent in error recovery for a single LBA can be limited using the Read Retry Count or Write Retry Count bytes in the Error Recovery mode page.

The drive firmware error recovery algorithms consist of 20 levels for read recoveries and six levels for write. Each level may consist of multiple steps, where a step is defined as a recovery function involving a single re-read or re-write attempt. The maximum level used by the drive in LBA recovery is determined by the read and write retry counts.

Table 5 equates the read and write retry count with the maximum possible recovery time for read and write recovery of individual LBAs. The times given do not include time taken to perform reallocations. Reallocations are performed when the ARRE bit (for reads) or AWRE bit (for writes) is one, the RC bit is zero, and the recovery time limit for the command has not yet been met. Time needed to perform reallocation is not counted against the recovery time limit.

When the RC bit is one, reallocations are disabled even if the ARRE or AWRE bits are one. The drive will still perform data recovery actions within the limits defined by the Read Retry Count, Write Retry Count, and Recovery Time Limit parameters. However, the drive does not report any unrecovered errors.

Table 5 Read and write retry count maximum recovery times

| Read retry count* | Maximum recovery time per LBA (cumulative, ms) | Write retry count | Maximum recovery time per LBA (cumulative, ms) |
|-------------------|---|-------------------|---|
| | | 0 | 25 |
| 1 | 90 | 1 | 50 |
| 5 | 410 | 2 | 78 |
| 10 | 830 | 3 | 95 |
| 15 | 1240 | 4 | 186 |
| 20 (default) | 1810 | 5 (default) | 295 |

* For read retry count, every tick ~ 5% of total error recovery. Valid range setting is 1-20.

e.g. 1 ~ 5%

5 ~ 25%

20 ~ 100%

Setting these retry counts to a value below the default setting could result in degradation of the unrecovered error rate. For example, suppose the read/write recovery page has the RC bit = 0 and if the read retry count is set to 5, this means ~ 25% of error recovery will be executed which consumes 621.62 ms (please refer to the table above). If the limit is reached and a LBA has not yet been recovered (i.e. requires retries beyond 621.62 ms), the command will end with Check Condition status report and unrecoverable read error will be reported.

9.3 SAS system errors

Information on the reporting of operational errors or faults across the interface is given in the *SAS Interface Manual*. The SSP Response returns information to the host about numerous kinds of errors or faults. The Receive Diagnostic Results reports the results of diagnostic operations performed by the drive.

Status returned by the drive to the initiator is described in the *SAS Interface Manual*. Status reporting plays a role in systems error management and its use in that respect is described in sections where the various commands are discussed.

9.4 Background Media Scan

Background Media Scan (BMS) is a self-initiated media scan. BMS is defined in the T10 document SPC-4 available from the T10 committee. BMS performs sequential reads across the entire pack of the media while the drive is idle. In RAID arrays, BMS allows hot spare drives to be scanned for defects prior to being put into service by the host system. On regular duty drives, if the host system makes use of the BMS Log Page, it can avoid placing data in suspect locations on the media. Unreadable and recovered error sites will be logged or reallocated per ARRE/AWRE settings.

With BMS, the host system can consume less power and system overhead by only checking BMS status and results rather than tying up the bus and consuming power in the process of host-initiated media scanning activity.

Since the background scan functions are only done during idle periods, BMS causes a negligible impact to system performance. The first BMS scan for a newly manufactured drive is performed as quickly as possible to verify the media and protect data by setting the “Start time after idle” to 5ms, all subsequent scans begin after 500ms of idle time. Other features that normally use idle time to function will function normally because BMS functions for limited time bursts then suspends activity to allow other background functions to operate.

BMS interrupts immediately to service host commands from the interface bus while performing reads. BMS will complete any BMS-initiated error recovery prior to returning to service host-initiated commands. Overhead associated with a return to host-servicing activity from BMS only impacts the first command that interrupted BMS, this results in a typical delay of about 1ms.

9.5 Deferred Auto-Reallocation

Deferred Auto-Reallocation (DAR) simplifies reallocation algorithms at the system level by allowing the drive to reallocate unreadable locations on a subsequent write command. Sites are marked for DAR during read operations performed by the drive. When a write command is received for an LBA marked for DAR, the auto-reallocation process is invoked and attempts to rewrite the data to the original location. If a verification of this rewrite fails, the sector is re-mapped to a spare location.

This is in contrast to the system having to use the Reassign Command to reassign a location that was unreadable and then generate a write command to rewrite the data. DAR is most effective when AWRE and ARRE are enabled—this is the default setting from the Seagate factory. With AWRE and ARRE disabled DAR is unable to reallocate the failing location and will report an error sense code indicating that a write command is being attempted to a previously failing location.

9.6 Idle Read After Write

Idle Read After Write (IRAW) utilizes idle time to verify the integrity of recently written data. During idle periods, no active system requests, the drive reads recently written data from the media and compares it to valid write command data resident in the drives data buffer. Any sectors that fail the comparison result in the invocation of a rewrite and auto-reallocation process. The process attempts to rewrite the data to the original location. If a verification of this rewrite fails, the sector is re-mapped to a spare location.

9.7 Protection Information (PI)

Protection Information is intended as a standardized approach to system level LRC traditionally provided by systems using 520 byte formatted LBAs. Drives formatted with PI information provide the same, common LBA count (i.e. same capacity point) as non-PI formatted drives. Sequential performance of a PI drive will be reduced by approximately 1.56% due to the extra overhead of PI being transferred from the media that is not calculated as part of the data transferred to the host. To determine the full transfer rate of a PI drive, transfers should be calculated by adding the 8 extra bytes of PI to the transferred LBA length, i.e. $512 + 8 = 520$. PI formatted drives are physically formatted to 520 byte sectors that store 512 bytes of customer data with 8 bytes of Protection Information appended to it. The advantage of PI is that the Protection Information bits can be managed at the HBA and HBA driver level. Allowing a system that typically does not support 520 LBA formats to integrate this level of protection.

Protection Information is valid with any supported LBA size. 512 LBA size is used here as common example.

9.7.1 Levels of PI

There are 4 types of Protection Information.

Type 0 - Describes a drive that is not formatted with PI information bytes. This allows for legacy support in non-PI systems.

Type 1 - Provides support of PI protection using 10 and 16 byte commands. The RDPROTECT and WRTPROTECT bits allow for checking control through the CDB. Eight bytes of Protection Information are transmitted at LBA boundaries across the interface if RDPROTECT and WRTPROTECT bits are nonzero values. Type 1 does not allow the use of 32 byte commands.

Type 2 - Provides checking control and additional expected fields within the 32 byte CDBs. Eight bytes of Protection Information are transmitted at LBA boundaries across the interface if RDPROTECT and WRTPROTECT bits are nonzero values. Type 2 does allow the use of 10 and 16 byte commands with zero values in the RDPROTECT and WRTPROTECT fields. The drive will generate 8 bytes (e.g.0xFFFF) 8 bytes of Protection Information to be stored on the media, but the 8 bytes will not be transferred to the host during a read command.

Type 3 - Seagate products do not support Type 3.

9.7.2 Setting and determining the current Type Level

A drive is initialized to a type of PI by using the format command on a PI capable drive. Once a drive is formatted to a PI Type, it may be queried by a Read Capacity (16) command to report the PI type which it is currently formatted to. PI Types cannot coexist on a single drive. A drive can only be formatted to a single PI Type. It can be changed at anytime to a new Type but requires a low level format which destroys all existing data on the drive. No other vehicle for changing the PI type is provided by the T10 SBC3 specification.

Type 1 PI format CDB command: 04 90 00 00 00 00, Write Buffer: 00 A0 00 00

Type 2 PI format CDB command: 04 D0 00 00 00 00, Write Buffer: 00 A0 00 00

9.7.3 Identifying a Protection Information drive

The Standard Inquiry provides a bit to indicate if PI is support by the drive. Vital Product Descriptor (VPD) page 0x86 provides bits to indicate the PI Types supported and which PI fields the drive supports checking.

NOTE

For further details with respect to PI, please refer to SCSI Block Commands - 3 (SBC-3) Draft Standard documentation.

9.8 Seagate RAID Rebuild™

Seagate RAID Rebuild is an industry standard feature to enable faster recovery from a failed drive in a RAID configuration. It improves RAID rebuild performance by extracting easily readable data from a failing drive. It quickly identifies blocks that would take longer to recover from the failed drive than to rebuild from the remaining drives in the RAID group.

This feature allows host control of error recovery, maximizes up time, and minimizes likelihood of 2nd drive failure in a RAID configuration.

10.0 Installation

Seagate Enterprise Performance 10K HDD disk drive installation is a plug-and-play process. There are no jumpers, switches, or terminators on the drive.

SAS drives are designed to be used in a host system that provides a SAS-compatible backplane with bays designed to accommodate the drive. In such systems, the host system typically provides a carrier or tray onto which the drive needs to be mounted. Mount the drive to the carrier or tray provided by the host system using four M3 x 0.5 metric screws. When tightening the screws, use a maximum torque of 4.5 in-lb +/- 0.45 in-lb. Do not over-tighten or force the screws. Mount the drive in any orientation.

NOTE

SAS drives are designed to be attached to the host system without I/O or power cables. If the intent is to use the drive in a non-backplane host system, connecting the drive using high-quality cables is acceptable as long as the I/O cable length does not exceed 10 meters (32.8 feet).

Slide the carrier or tray into the appropriate bay in the host system using the instructions provided by the host system. This connects the drive directly to the system's SAS connector. The SAS connector is normally located on a SAS backpanel. See Section 11.4.1 for additional information about these connectors.

Power is supplied through the SAS connector.

The drive is shipped from the factory low-level formatted in 512-byte logical blocks for 512 emulation and native drives and 4096-byte logical blocks for 4096 native drives. Reformat the drive only if a different logical block size needs to be selected.




Figure 12. Physical interface

NOTE

Image is for reference only, may not represent actual drive.

10.1 Drive orientation

The drive may be mounted in any orientation. All drive performance characterizations, however, have been done with the drive in horizontal (discs level) and vertical (drive on its side) orientations, which are the two preferred mounting orientations.

10.2 Cooling

Cabinet cooling must be designed by the customer so that the ambient temperature immediately surrounding the drive does not result in the drive exceeding temperature conditions specified in Section 6.5.1, "Temperature".

The rack, cabinet, or drawer environment for the drive must provide heat removal from the electronics and head and disk assembly (HDA). Confirm that adequate heat removal is provided using the temperature measurement guidelines described in Section 6.5.1.

Forced air flow may be required to keep temperatures at or below the temperatures specified in Section 6.5.1 in which case the drive should be oriented, or air flow directed, so that the least amount of air flow resistance is created while providing air flow to the electronics and HDA. Also, the shortest possible path between the air inlet and exit should be chosen to minimize the travel length of air heated by the drive and other heat sources within the rack, cabinet, or drawer environment.

If forced air is determined to be necessary, possible air-flow patterns are shown in Figure 13. The air-flow patterns are created by one or more fans, either forcing or drawing air as shown in the illustrations. Conduction, convection, or other forced air-flow patterns are acceptable as long as the temperature measurement guidelines of Section 6.5.1 are met.




Figure 13. Air flow

NOTE Image of the HDA may not represent actual product, for reference only.

10.3 Drive mounting

Mount the drive using the bottom or side mounting holes. If mounting the drive using the bottom holes, ensure to not physically distort the drive by attempting to mount it on a stiff, non-flat surface.

The allowable mounting surface stiffness is 80 lb/in (14.0 N/mm). The following equation and paragraph define the allowable mounting surface stiffness:

$$K \times X = F < 15\text{lb} = 67\text{N}$$

where K is the mounting surface stiffness (units in lb/in or N/mm) and X is the out-of-plane surface distortion (units in inches or millimeters). The out-of-plane distortion (X) is determined by defining a plane with three of the four mounting points fixed and evaluating the out-of-plane deflection of the fourth mounting point when a known force (F) is applied to the fourth point.

10.4 Grounding

Signal ground (PCBA) and HDA ground are connected together in the drive and cannot be separated by the user. The equipment in which the drive is mounted is connected directly to the HDA and PCBA with no electrically isolating shock mounts. If it is desired for the system chassis to not be connected to the HDA/PCBA ground, the systems integrator or user must provide a nonconductive (electrically isolating) method of mounting the drive in the host equipment.

Increased radiated emissions may result if maximum surface area ground connection is not provided between system ground and drive ground. This is the system designer's and integrator's responsibility.

11.0 Interface requirements

This section partially describes the interface requirements as implemented on Seagate Enterprise Performance 10K HDD drives. Additional information is provided in the *SAS Interface Manual* (part number 100293071).

11.1 SAS features

This section lists the SAS-specific features supported by Seagate Enterprise Performance 10K HDD drives.

11.1.1 Task management functions

The table below lists the SAS task management functions supported.

Table 6 *SAS task management functions supported*

| Task name | Supported |
|--------------------|-----------|
| Abort Task | Yes |
| Clear ACA | Yes |
| Clear task set | Yes |
| Abort task set | Yes |
| Logical Unit Reset | Yes |
| Query Task | Yes |

11.1.2 Task management responses

The table below lists the SAS response codes returned for task management functions supported.

Table 7 *Task management response codes*

| Function name | Response code |
|------------------------|---------------|
| Function complete | 00 |
| Invalid frame | 02 |
| Function not supported | 04 |
| Function failed | 05 |
| Function succeeded | 08 |
| Invalid logical unit | 09 |

11.2 Dual port support

Seagate Enterprise Performance 10K HDD SAS drives have two independent ports. These ports may be connected in the same or different SCSI domains. Each drive port has a unique SAS address.

The two ports have the capability of independent port clocking (e.g. both ports can run at 12Gb/s or the first port can run at 12Gb/s while the second port runs at 3Gb/s.) The supported link rates are 3.0, 6.0 or 12.0 Gb/s.

Subject to buffer availability, the Seagate Enterprise Performance 10K HDD drives support:

- Concurrent port transfers—The drive supports receiving COMMAND, TASK management transfers on both ports at the same time.
- Full duplex—The drive supports sending XFER_RDY, DATA and RESPONSE transfers while receiving frames on both ports.

11.3 SCSI commands supported

The table below lists the SCSI commands supported by Seagate Enterprise Performance 10K HDD drives.

Table 8 Supported commands

| Command name | Command code | Supported |
|--|--------------|-----------|
| Change Definition | 40h | N |
| Compare | 39h | N |
| Copy | 18h | N |
| Copy and Verify | 3Ah | N |
| Format Unit ^{[1][5]} | 04h | Y |
| DCRT bit supported | | Y |
| DPRY bit supported | | N |
| DSP bit supported | | Y |
| IMMED bit supported | | Y |
| IP bit supported | | Y |
| SI (Security Initialize) bit supported | | N |
| STPF bit supported | | Y |
| VS (vendor specific) | | N |
| Inquiry | 12h | Y |
| Block Device Characteristics page (B1h) | | Y |
| Block Limits page (B0h) | | Y |
| Date Code page (C1h) | | Y |
| Device Behavior page (C3h) | | Y |
| Device Identification page (83h) | | Y |
| Extended Inquiry Data page (86h) | | Y |
| Firmware Numbers page (C0h) | | Y |
| Jumper Settings page (C2h) | | Y |
| Mode Page policy page (87h) | | N |
| Power Condition page (8Ah) | | Y |
| Power Consumption page (8Dh) | | N |
| Protocol Specific Logical unit Information page (90h) | | N |
| Protocol Specific Port Information page (91h) | | N |
| SCSI Ports page (88h) | | Y |
| Supported Vital Product Data page (00h) | | Y |
| Unit Serial Number page (80h) | | Y |
| Lock-unlock cache | 36h | N |
| Log Select (same log pages as Log Sense 4Dh) | 4Ch | Y |
| PCR bit | | Y |
| DU bit | | N |
| DS bit | | Y |
| TSD bit | | Y |
| ETC bit | | N |
| TMC bit | | N |
| LP bit | | N |
| Log Sense | 4Dh | Y |
| Application Client Log page (0Fh) | | Y |
| Buffer Over-run/Under-run page (01h) | | N |
| Cache Statistics page (37h) | | Y |
| Factory Log page (3Eh) | | Y |
| Information Exceptions Log page (2Fh) | | Y |
| Last n Deferred Errors or Asynchronous Events page (0Bh) | | N |

Table 8 Supported commands

| Command name | Command code | Supported |
|--|--------------|-------------------------|
| Last n Error Events page (07h) | | N |
| Non-medium Error page (06h) | | Y |
| Pages Supported list (00h) | | Y |
| Protocol-specific Log Page for SAS (18h) | | Y |
| Read Error Counter page (03h) | | Y |
| Read Reverse Error Counter page (04h) | | N |
| Self-test Results page (10h) | | Y |
| Start-stop Cycle Counter page (0Eh) | | Y |
| Temperature page (0Dh) | | Y |
| Verify Error Counter page (05h) | | Y |
| Write error counter page (02h) | | Y |
| Mode Select (same pages as Mode Sense 1Ah) ^[3] | 15h | Y ^[2] |
| Mode Select (10) (same pages as Mode Sense 1Ah) ^[3] | 55h | Y |
| Mode Sense | 1Ah | Y ^[2] |
| Caching Parameters page (08h) | | Y |
| Control Mode page (0Ah) | | Y |
| Control Extension Mode Page (01h) | | Y |
| Disconnect/Reconnect (02h) | | Y |
| Error Recovery page (01h) | | Y |
| Format page (03h) | | Y |
| Information Exceptions Control page (1Ch) | | Y |
| Background Scan mode subpage (01h) | | Y |
| Notch and Partition Page (0Ch) | | N |
| Protocol-Specific Port page (19h) | | Y |
| Power Condition page (1Ah) | | Y |
| Rigid Disk Drive Geometry page (04h) | | Y |
| Unit Attention page (00h) | | Y |
| Verify Error Recovery page (07h) | | Y |
| Xor Control page (10h) | | N |
| Mode Sense (10) (same pages as Mode Sense 1Ah) ^[3] | 5Ah | Y |
| Persistent Reserve In | 5Eh | Y |
| Persistent Reserve Out | 5Fh | Y |
| Prefetch | 34h | N |
| Read (6) | 08h | Y |
| Read (10) | 28h | Y |
| DPO bit supported | | Y |
| FUA bit supported | | Y |
| Read (12) | A8h | N |
| Read (16) | 88h | Y |
| Read (32) | 7Fh/0009h | Y (Type 2 PI only) |
| Read Buffer (modes 0, 2, 3, Ah and Bh supported) | 3Ch | Y (non-SED drives only) |
| Read Capacity (10) | 25h | Y |
| Read Capacity (16) | 9Eh/10h | Y |
| Read Defect Data (10) | 37h | Y |
| Read Defect Data (12) | B7h | Y |
| Read Long | 3Eh | Y (non-SED drives only) |
| Read Long (16) | 9Eh/11h | Y |
| Reassign Blocks | 07h | Y |
| Receive Diagnostic Results | 1Ch | Y |

Table 8 Supported commands

| Command name | Command code | Supported |
|--|--------------|---------------------|
| Supported Diagnostics pages (00h) | | Y |
| Translate page (40h) | | Y |
| Release | 17h | Y |
| Release (10) | 57h | Y |
| Report LUNs | A0h | Y |
| Report Timestamp | A3h | N |
| Request Sense | 03h | Y |
| Actual Retry Count bytes | | Y |
| Extended Sense | | Y |
| Field Pointer bytes | | Y |
| Reserve | 16h | Y |
| 3rd Party Reserve | | Y |
| Extent Reservation | | N |
| Reserve (10) | 56h | Y |
| 3rd Party Reserve | | Y |
| Extent Reservation | | N |
| Rezero Unit | 01h | Y |
| Sanitize | 48h | Y |
| Block Erase | | N |
| Cryptographic Erase | | Y (SED only) |
| Overwrite | | N |
| Sanitize Exit | | N |
| Security Protocol In | A2h | Y (SED models only) |
| Security Protocol Out | B5h | Y (SED models only) |
| Seek (6) | 0Bh | Y |
| Seek (10) | 2Bh | Y |
| Send Diagnostics | 1Dh | Y |
| Rebuild Assist page (42h) | | Y |
| Supported Diagnostics pages (00h) | | Y |
| Translate page (40h) | | Y |
| Set Timestamp | 0Fh | Y |
| Start Unit/Stop Unit (spindle ceases rotating) | 1Bh | Y |
| Synchronize Cache | 35h | Y |
| Synchronize Cache (16) | 91h | Y |
| Test Unit Ready | 00h | Y |
| Verify (10) | 2Fh | Y |
| BYTCHK bit | | Y |
| Verify (12) | AFh | N |
| Verify (16) | AFh | Y |
| Verify (32) | 7Fh/000Ah | Y (Type 2 PI only) |
| Write (6) | 0Ah | Y |
| Write (10) | 2Ah | Y |
| DPO bit | | Y |
| FUA bit | | Y |
| Write (12) | AAh | N |
| Write (16) | 8Ah | Y |
| Write (32) | 7Fh/000Bh | Y (Type 2 PI only) |
| Write and Verify (10) | 2Eh | Y |
| DPO bit | | Y |

Table 8 Supported commands

| Command name | Command code | Supported |
|--|--------------|-------------------------|
| Write and Verify (12) | AEh | N |
| Write and Verify (16) | 8Eh | Y |
| Write and Verify (32) | 7Fh/000Ch | Y (Type 2 PI only) |
| Write Buffer (modes 0, 2, supported) | 3Bh | Y (non-SED drives only) |
| Write Buffer | 3Bh | |
| Firmware Download option (modes 5, 7, Dh and Fh) | | Y (non-SED drives only) |
| Firmware Download option (modes 4, 5, 7) | | Y (SED drives only) |
| Write Long (10) | 3Fh | Y |
| Write Long (16) | 9Fh/11h | Y |
| Write Same (10) ^[5] | 41h | Y |
| PBdata | | N |
| LBdata | | N |
| Write Same (16) ^[5] | 93h | Y |
| Write Same (32) | 7Fh/000Dh | Y (Type 2 PI only) |
| XDRead | 52h | N |
| XDWrite | 50h | N |
| XPWrite | 51h | N |

- [1] Seagate Enterprise Performance 10K HDD drives can format to 512, 520, 524, 528 (512 emulation) or 4096, 4160, 4192, or 4224 (4096 native).
- [2] **Warning.** Power loss during flash programming can result in firmware corruption. This usually makes the drive inoperable.
- [3] Reference Mode Sense command 1Ah for mode pages supported.
- [4] Y = Yes. Command is supported.
N = No. Command is not supported.
A = Support is available on special request.
- [5] Approximately 1.5 increase in time to complete this command for a SED drive versus a non-SED drive of the same capacity.

11.3.1 Inquiry data

The table below lists the Inquiry command data that the drive should return to the initiator per the format given in the *SAS Interface Manual*.

Table 9 Seagate Enterprise Performance 10K HDD inquiry data

| Bytes | Data (hex) | | | | | | | | | | | | | | | | |
|---------|------------|-----|------|----|----|----|----|----|----|----|----|-----|----|----|----|-----|------------|
| 0-15 | 00 | 00 | xx** | 12 | 8B | 01 | PP | 02 | 53 | 45 | 41 | 47 | 41 | 54 | 45 | 20 | Vendor ID |
| 16-31 | [53 | 54 | 31 | 38 | 30 | 30 | 4D | 4D | 30 | 30 | 30 | 38] | 20 | 20 | 20 | 20 | Product ID |
| 32-47 | R# | R# | R# | R# | S# | S# | S# | S# | S# | S# | S# | S# | 00 | 00 | 00 | 00 | |
| 48-63 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | |
| 64-79 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | |
| 80-95 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | |
| 96-111 | 00 | 43 | 6F | 70 | 79 | 72 | 69 | 67 | 68 | 74 | 20 | 28 | 63 | 29 | 20 | 32* | *Copyright |
| 112-127 | 30* | 31* | 33* | 20 | 53 | 65 | 61 | 67 | 61 | 74 | 65 | 20 | 41 | 6C | 6C | 20 | notice |
| 128-143 | 72 | 69 | 67 | 68 | 74 | 73 | 20 | 72 | 65 | 73 | 65 | 72 | 76 | 65 | 64 | 20 | |

* Copyright year (changes with actual year).

** SCSI Revision support. See the appropriate SPC release documentation for definitions.

PP 10 = Inquiry data for an Inquiry command received on Port A.

30 = Inquiry data for an Inquiry command received on Port B.

R# Four ASCII digits representing the last four digits of the product firmware release number.

S# Eight ASCII digits representing the eight digits of the product serial number.

[] Bytes 16 through 26 reflect model of drive. The table above shows the hex values for Model ST1800MM0008.

Refer to the values below for the values of bytes 16 through 26 of a particular model:

| | | | | | | | | | | | | |
|--------------|----|----|----|----|----|----|----|----|----|----|----|----|
| ST1800MM0018 | 53 | 54 | 31 | 38 | 30 | 30 | 4D | 4D | 30 | 30 | 31 | 38 |
| ST1800MM0038 | 53 | 54 | 31 | 38 | 30 | 30 | 4D | 4D | 30 | 30 | 33 | 38 |
| ST1800MM0048 | 53 | 54 | 31 | 38 | 30 | 30 | 4D | 4D | 30 | 30 | 34 | 38 |
| ST1800MM0068 | 53 | 54 | 31 | 38 | 30 | 30 | 4D | 4D | 30 | 30 | 36 | 38 |
| ST1800MM0078 | 53 | 54 | 31 | 38 | 30 | 30 | 4D | 4D | 30 | 30 | 37 | 38 |
| ST1800MM0088 | 53 | 54 | 31 | 38 | 30 | 30 | 4D | 4D | 30 | 30 | 38 | 38 |
| ST1800MM0108 | 53 | 54 | 31 | 38 | 30 | 30 | 4D | 4D | 30 | 31 | 30 | 38 |
| ST1800MM0118 | 53 | 54 | 31 | 38 | 30 | 30 | 4D | 4D | 30 | 31 | 31 | 38 |
| ST1800MM0128 | 53 | 54 | 31 | 38 | 30 | 30 | 4D | 4D | 30 | 31 | 32 | 38 |
| ST1800MM0148 | 53 | 54 | 31 | 38 | 30 | 30 | 4D | 4D | 30 | 31 | 34 | 38 |
| ST1800MM0158 | 53 | 54 | 31 | 38 | 30 | 30 | 4D | 4D | 30 | 31 | 35 | 38 |
| ST1200MM0008 | 53 | 54 | 31 | 32 | 30 | 30 | 4D | 4D | 30 | 30 | 30 | 38 |
| ST1200MM0018 | 53 | 54 | 31 | 32 | 30 | 30 | 4D | 4D | 30 | 30 | 31 | 38 |
| ST1200MM0038 | 53 | 54 | 31 | 32 | 30 | 30 | 4D | 4D | 30 | 30 | 33 | 38 |
| ST1200MM0068 | 53 | 54 | 31 | 32 | 30 | 30 | 4D | 4D | 30 | 30 | 36 | 38 |
| ST1200MM0088 | 53 | 54 | 31 | 32 | 30 | 30 | 4D | 4D | 30 | 30 | 38 | 38 |
| ST1200MM0098 | 53 | 54 | 31 | 32 | 30 | 30 | 4D | 4D | 30 | 30 | 39 | 38 |
| ST1200MM0118 | 53 | 54 | 31 | 32 | 30 | 30 | 4D | 4D | 30 | 31 | 31 | 38 |
| ST1200MM0138 | 53 | 54 | 31 | 32 | 30 | 30 | 4D | 4D | 30 | 31 | 33 | 38 |
| ST1200MM0158 | 53 | 54 | 31 | 32 | 30 | 30 | 4D | 4D | 30 | 31 | 35 | 38 |
| ST1200MM0178 | 53 | 54 | 31 | 32 | 30 | 30 | 4D | 4D | 30 | 31 | 37 | 38 |
| ST900MM0008 | 53 | 54 | 39 | 30 | 30 | 30 | 4D | 4D | 30 | 30 | 30 | 38 |
| ST900MM0018 | 53 | 54 | 39 | 30 | 30 | 30 | 4D | 4D | 30 | 30 | 31 | 38 |
| ST900MM0038 | 53 | 54 | 39 | 30 | 30 | 30 | 4D | 4D | 30 | 30 | 33 | 38 |
| ST900MM0068 | 53 | 54 | 39 | 30 | 30 | 30 | 4D | 4D | 30 | 30 | 36 | 38 |
| ST900MM0088 | 53 | 54 | 39 | 30 | 30 | 30 | 4D | 4D | 30 | 30 | 38 | 38 |
| ST900MM0108 | 53 | 54 | 39 | 30 | 30 | 30 | 4D | 4D | 30 | 31 | 30 | 38 |
| ST900MM0128 | 53 | 54 | 39 | 30 | 30 | 30 | 4D | 4D | 30 | 31 | 32 | 38 |
| ST900MM0148 | 53 | 54 | 39 | 30 | 30 | 30 | 4D | 4D | 30 | 31 | 34 | 38 |
| ST600MM0008 | 53 | 54 | 36 | 30 | 30 | 30 | 4D | 4D | 30 | 30 | 30 | 38 |
| ST600MM0018 | 53 | 54 | 36 | 30 | 30 | 30 | 4D | 4D | 30 | 30 | 31 | 38 |
| ST600MM0038 | 53 | 54 | 36 | 30 | 30 | 30 | 4D | 4D | 30 | 30 | 33 | 38 |
| ST600MM0068 | 53 | 54 | 36 | 30 | 30 | 30 | 4D | 4D | 30 | 30 | 36 | 38 |

| | | | | | | | | | | | |
|-------------|----|----|----|----|----|----|----|----|----|----|----|
| ST600MM0088 | 53 | 54 | 36 | 30 | 30 | 4D | 4D | 30 | 30 | 38 | 38 |
| ST600MM0098 | 53 | 54 | 36 | 30 | 30 | 4D | 4D | 30 | 30 | 39 | 38 |
| ST600MM0118 | 53 | 54 | 36 | 30 | 30 | 4D | 4D | 30 | 31 | 31 | 38 |
| ST600MM0138 | 53 | 54 | 36 | 30 | 30 | 4D | 4D | 30 | 31 | 33 | 38 |
| ST600MM0158 | 53 | 54 | 36 | 30 | 30 | 4D | 4D | 30 | 31 | 35 | 38 |
| ST600MM0178 | 53 | 54 | 36 | 30 | 30 | 4D | 4D | 30 | 31 | 37 | 38 |

11.3.2 Mode Sense data

The Mode Sense command provides a way for the drive to report its operating parameters to the initiator. The drive maintains four sets of mode parameters:

1. Default values

Default values are hard-coded in the drive firmware stored in flash E-PROM (nonvolatile memory) on the drive's PCB. These default values can be changed only by downloading a complete set of new firmware into the flash E-PROM. An initiator can request and receive from the drive a list of default values and use those in a Mode Select command to set up new current and saved values, where the values are changeable.

2. Saved values

Saved values are stored on the drive's media using a Mode Select command. Only parameter values that are allowed to be changed can be changed by this method. Parameters in the saved values list that are not changeable by the Mode Select command get their values from default values storage.

When power is applied to the drive, it takes saved values from the media and stores them as current values in volatile memory. It is not possible to change the current values (or the saved values) with a Mode Select command before the drive achieves operating speed and is "ready." An attempt to do so results in a "Check Condition" status.

On drives requiring unique saved values, the required unique saved values are stored into the saved values storage location on the media prior to shipping the drive. Some drives may have unique firmware with unique default values also.

On standard OEM drives, the saved values are taken from the default values list and stored into the saved values storage location on the media prior to shipping.

3. Current values

Current values are volatile values being used by the drive to control its operation. A Mode Select command can be used to change the values identified as changeable values. Originally, current values are installed from saved or default values after a power on reset, hard reset, or Bus Device Reset message.

4. Changeable values

Changeable values form a bit mask, stored in nonvolatile memory, that dictates which of the current values and saved values can be changed by a Mode Select command. A one (1) indicates the value can be changed. A zero (0) indicates the value is not changeable. For example, in the table in Section 11.3.2, refer to Mode page 81, in the row entitled "CHG." These are hex numbers representing the changeable values for Mode page 81. Note in columns 5 and 6 (bytes 04 and 05), there is 00h which indicates that in bytes 04 and 05 none of the bits are changeable. Note also that bytes 06, 07, 09, 10, and 11 are not changeable, because those fields are all zeros. In byte 02, hex value FF equates to the binary pattern 11111111. If there is a zero in any bit position in the field, it means that bit is not changeable. Since all of the bits in byte 02 are ones, all of these bits are changeable.

The changeable values list can only be changed by downloading new firmware into the flash E-PROM.

NOTE

Because there are often several different versions of drive control firmware in the total population of drives in the field, the Mode Sense values given in the following tables may not exactly match those of some drives.

The following tables list the values of the data bytes returned by the drive in response to the Mode Sense command pages for SCSI implementation (see the *SAS Interface Manual*).

DEF = Default value. Standard OEM drives are shipped configured this way.

CHG = Changeable bits; indicates if default value is changeable.

Table 10 Mode Sense data for 1.8TB drives

BLOCK DESCRIPTOR:

00 00 00 00 1A 32 DD 36 00 00 10 00 (4K native)
00 00 00 00 D1 96 E9 B0 00 00 02 00 (5xx emulation)

DEF: 01 0A C0 14 FF 00 00 00 05 00 FF FF
CHG: 01 0A FF FF 00 00 00 00 FF 00 FF FF

DEF: 02 0E 00 00 00 00 00 00 00 00 A0 00 00 00 00
CHG: 02 0E 00 00 FF FF 00 00 FF FF FF FF 00 00 00 00

DEF: 03 16 BB D0 00 00 00 00 03 80 04 C4 10 00 00 01 00 8E 00 18 40 00 00 00
CHG: 03 16 00

DEF: 04 16 03 43 28 08 00 00 00 00 00 00 00 00 00 00 00 00 29 04 00 00
CHG: 04 16 00

DEF: 07 0A 00 14 FF 00 00 00 00 00 FF FF
CHG: 07 0A 0F FF 00 00 00 00 00 00 FF FF

DEF: 08 12 14 00 FF FF 00 00 FF FF FF FF 90 20 00 00 00 00 00
DEF: 08 12 10 00 FF FF 00 00 FF FF FF FF 90 20 00 00 00 00 00 (TurboBoost models)
CHG: 08 12 A5 00 00 00 FF FF FF FF 00 00 30 00 00 00 00 00 00 00

DEF: 0A 0A 02 00 00 80 00 00 00 00 2C 56
CHG: 0A 0A 0F F6 00 00 00 00 00 00 00 00

DEF: 18 06 06 00 00 00 00 00
CHG: 18 06 00 00 00 00 00 00

DEF: 19 0E 46 00 07 D0 00 00 00 00 00 00 00 00 00
CHG: 19 0E 50 00 FF FF FF FF FF FF 00 00 00 00 00 00

DEF: 1A 26 00 06 00 00 00 0A 00 00 8C A0 00 00 17 70 00 00 46 50 00 00 46 50 00 00 00 00 00 00 00 00 00 00 00 00 58
CHG: 1A 26 01 0F FF 00 00 00 00 00 00 00 00 00 00 00 00 00 FC

DEF: 1C 0A 10 00 00 00 00 00 00 00 00 01
CHG: 1C 0A 9D 0F FF FF FF FF FF FF FF FF

DEF: 00 0A 00 80 0F 00 00 00 00 00 00 00
CHG: 00 0A B7 C0 8F 00 00 00 00 00 00 FF FF

Table 11 Mode Sense data for 1.2TB drives

BLOCK DESCRIPTOR:

00 00 00 00 11 77 41 96 00 00 10 00 (4K native)

00 00 00 00 8B BA 0C B0 00 00 02 00 (5xx emulation and 5xx native)

DEF: 01 0A C0 14 FF 00 00 00 05 00 FF FF

CHG: 01 0A FF FF 00 00 00 00 FF 00 FF FF

DEF: 02 0E 00 00 00 00 00 00 00 00 00 A0 00 00 00 00

CHG: 02 0E 00 00 FF FF 00 00 FF FF FF FF 00 00 00 00

DEF: 03 16 BB D0 00 00 00 00 03 80 04 C4 10 00 00 01 00 8E 00 18 40 00 00 00 (4K native and 5xx emulation)

DEF: 03 16 BB D0 00 00 00 00 03 80 04 C4 02 00 00 01 00 8E 00 18 40 00 00 00 (5xx native)

CHG: 03 16 00

DEF: 04 16 03 43 28 06 00 00 00 00 00 00 00 00 00 00 00 00 29 04 00 00

CHG: 04 16 00

DEF: 07 0A 00 14 FF 00 00 00 00 00 FF FF

CHG: 07 0A 0F FF 00 00 00 00 00 00 FF FF

DEF: 08 12 14 00 FF FF 00 00 FF FF FF FF 90 20 00 00 00 00 00 00 (4K native and 5xx emulation)

DEF: 08 12 10 00 FF FF 00 00 FF FF FF FF 90 20 00 00 00 00 00 00 (TurboBoost models)

CHG: 08 12 A5 00 00 00 FF FF FF FF 00 00 30 00 00 00 00 00 00 (4K native and 5xx emulation)

DEF: 08 12 14 00 FF FF 00 00 FF FF FF FF 80 20 00 00 00 00 00 00 (5xx native)

CHG: 08 12 A5 00 00 00 FF FF FF FF 00 00 20 00 00 00 00 00 00 (5xx native)

DEF: 0A 0A 02 00 00 80 00 00 00 00 1D CE

CHG: 0A 0A 0F F6 00 00 00 00 00 00 00 00

DEF: 18 06 06 00 00 00 00 00

CHG: 18 06 00 00 00 00 00 00

DEF: 19 0E 46 00 07 D0 00 00 00 00 00 00 00 00 00 00

CHG: 19 0E 50 00 FF FF FF FF FF FF 00 00 00 00 00 00

DEF: 1A 26 00 06 00 00 00 0A 00 00 8C A0 00 00 17 70 00 00 46 50 00 00 46 50 00 00 00 00 00 00 00 00 00 00 00 58

CHG: 1A 26 01 0F FF 00 00 00 00 00 00 00 00 00 00 00 00 FC

DEF: 1C 0A 10 00 00 00 00 00 00 00 00 01

CHG: 1C 0A 9D 0F FF FF FF FF FF FF FF FF FF

DEF: 00 0A 00 80 0F 00 00 00 00 00 00 00

CHG: 00 0A B7 C0 8F 00 00 00 00 00 00 FF FF

Table 12 Mode Sense data for 900GB drives

BLOCK DESCRIPTOR:

00 00 00 00 0D 19 73 C6 00 00 10 00 (4K native)

00 00 00 00 68 CB 9E 30 00 00 02 00 (5xx emulation)

DEF: 01 0A C0 14 FF 00 00 00 05 00 FF FF

CHG: 01 0A FF FF 00 00 00 00 FF 00 FF FF

DEF: 02 0E 00 00 00 00 00 00 00 00 00 A0 00 00 00 00

CHG: 02 0E 00 00 FF FF 00 00 FF FF FF FF 00 00 00 00

DEF: 03 16 BB D0 00 00 00 00 03 80 04 C4 10 00 00 01 00 96 00 18 40 00 00 00

CHG: 03 16 00

DEF: 04 16 03 43 28 04 00 00 00 00 00 00 00 00 00 00 00 00 29 04 00 00

CHG: 04 16 00

DEF: 07 0A 00 14 FF 00 00 00 00 00 FF FF

CHG: 07 0A 0F FF 00 00 00 00 00 00 FF FF

DEF: 08 12 14 00 FF FF 00 00 FF FF FF FF 90 20 00 00 00 00 00

DEF: 08 12 10 00 FF FF 00 00 FF FF FF FF 90 20 00 00 00 00 00 (TurboBoost models)

CHG: 08 12 A5 00 00 00 FF FF FF FF 00 00 30 00 00 00 00 00 00 00

DEF: 0A 0A 02 00 00 80 00 00 00 00 16 3A

CHG: 0A 0A 0F F6 00 00 00 00 00 00 00 00

DEF: 18 06 06 00 00 00 00 00

CHG: 18 06 00 00 00 00 00 00

DEF: 19 0E 46 00 07 D0 00 00 00 00 00 00 00 00 00 00

CHG: 19 0E 50 00 FF FF FF FF FF FF 00 00 00 00 00 00

DEF: 1A 26 00 06 00 00 00 0A 00 00 8C A0 00 00 17 70 00 00 46 50 00 00 46 50 00 00 00 00 00 00 00 00 00 00 00 00 58

CHG: 1A 26 01 0F FF 00 00 00 00 00 00 00 00 00 00 00 00 FC

DEF: 1C 0A 10 00 00 00 00 00 00 00 00 01

CHG: 1C 0A 9D 0F FF FF FF FF FF FF FF FF FF

DEF: 00 0A 00 80 0F 00 00 00 00 00 00 00

CHG: 00 0A B7 C0 8F 00 00 00 00 00 00 FF FF

Table 13 Mode Sense data for 600GB drives

BLOCK DESCRIPTOR:

00 00 00 00 08 BB A5 F6 00 00 10 00 (4K native)

00 00 00 00 45 DD 2F B0 00 00 02 00 (5xx emulation and 5xx native)

DEF: 01 0A C0 14 FF 00 00 00 05 00 FF FF

CHG: 01 0A FF FF 00 00 00 00 FF 00 FF FF

DEF: 02 0E 00 00 00 00 00 00 00 00 00 A0 00 00 00 00

CHG: 02 0E 00 00 FF FF 00 00 FF FF FF FF 00 00 00 00

DEF: 03 16 BB D0 00 00 00 00 03 80 04 C4 10 00 00 01 00 96 00 18 40 00 00 00 (4K native and 5xx emulation)

DEF: 03 16 BB D0 00 00 00 00 03 80 04 C4 02 00 00 01 00 8E 00 18 40 00 00 00 (5xx native)

CHG: 03 16 00

DEF: 04 16 03 43 28 03 00 00 00 00 00 00 00 00 00 00 00 00 29 04 00 00

CHG: 04 16 00

DEF: 07 0A 00 14 FF 00 00 00 00 00 FF FF

CHG: 07 0A 0F FF 00 00 00 00 00 00 FF FF

DEF: 08 12 14 00 FF FF 00 00 FF FF FF FF 90 20 00 00 00 00 00 00 (4K native and 5xx emulation)

DEF: 08 12 10 00 FF FF 00 00 FF FF FF FF 90 20 00 00 00 00 00 00 (TurboBoost models)

CHG: 08 12 A5 00 00 00 FF FF FF FF 00 00 30 00 00 00 00 00 00 (4K native and 5xx emulation)

DEF: 08 12 14 00 FF FF 00 00 FF FF FF FF 80 20 00 00 00 00 00 00 (5xx native)

CHG: 08 12 A5 00 00 00 FF FF FF FF 00 00 20 00 00 00 00 00 00 (5xx native)

DEF: 0A 0A 02 00 00 80 00 00 00 00 0F 00

CHG: 0A 0A 0F F6 00 00 00 00 00 00 00 00

DEF: 18 06 06 00 00 00 00 00

CHG: 18 06 00 00 00 00 00 00

DEF: 19 0E 46 00 07 D0 00 00 00 00 00 00 00 00 00 00

CHG: 19 0E 50 00 FF FF FF FF FF FF 00 00 00 00 00 00

DEF: 1A 26 00 06 00 00 00 0A 00 00 8C A0 00 00 17 70 00 00 46 50 00 00 46 50 00 00 00 00 00 00 00 00 00 00 00 00 58

CHG: 1A 26 01 0F FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF FF 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 FC

DEF: 1C 0A 10 00 00 00 00 00 00 00 00 01

CHG: 1C 0A 9D 0F FF FF FF FF FF FF FF FF

DEF: 00 0A 00 80 0F 00 00 00 00 00 00 00

CHG: 00 0A B7 C0 8F 00 00 00 00 00 FF FF

11.4 Miscellaneous operating features and conditions

The table below lists various features and conditions. A “Y” in the support column indicates the feature or condition is supported. An “N” in the support column indicates the feature or condition is not supported.

Table 14 *Miscellaneous features*

| Supported | Feature or condition |
|-----------|---|
| N | Automatic contingent allegiance |
| N | Asynchronous event notification |
| N | Synchronized (locked) spindle operation |
| Y | Segmented caching |
| N | Zero latency read |
| Y | Queue tagging (up to 128 queue tags supported) |
| Y | Deferred error handling |
| Y | Parameter rounding (controlled by Round bit in Mode Select page 0) |
| Y | Reporting actual retry count in Extended Sense bytes 15, 16, and 17 |
| N | Adaptive caching |
| Y | SMP = 1 in Mode Select command needed to save RPL and rotational offset bytes |

Table 15 *Miscellaneous status*

| Supported | Status |
|-----------|-------------------------------|
| Y | Good |
| Y | Check condition |
| Y | Condition met/good |
| Y | Busy |
| Y | Reservation conflict |
| Y | Task set full |
| N | ACA active |
| N | ACA active, faulted initiator |

11.4.1 SAS physical interface

Figure 14 shows the location of the SAS device connector J1. Figures 15 and 16 provide the dimensions of the SAS connector.

Details of the physical, electrical, and logical characteristics are provided within this section. The operational aspects of Seagate’s SAS drives are provided in the *SAS Interface Manual*



Figure 14. Physical interface




Figure 16. SAS device plug dimensions (detail)

11.4.2 Physical characteristics

This section defines physical interface connector.

11.4.3 Connector requirements

Contact a preferred connector manufacturer for mating part information. Part numbers for SAS connectors will be provided in a future revision of this publication when production parts are available from major connector manufacturers.

The SAS device connector is illustrated in Figures 15 and 16.

11.4.4 Electrical description

SAS drives use the device connector for:

- DC power
- SAS interface
- Activity LED

This connector is designed to either plug directly into a backpanel or accept cables.

11.4.5 Pin descriptions

This section provides a pin-out of the SAS device and a description of the functions provided by the pins.

Table 16 SAS pin descriptions

| Pin | Signal name | Signal type | Pin | Signal name | Signal type |
|------|---------------|------------------|------|------------------------|--------------------|
| S1 | Port A Ground | | P1* | NC (reserved 3.3Volts) | |
| S2* | +Port A_in | Diff. input pair | P2* | NC (reserved 3.3Volts) | |
| S3* | -Port A_in | | P3 | SAS Power Disable ** | |
| S4 | Port A Ground | | P4 | Ground | |
| S5* | -Port A_out | Diff output pair | P5 | Ground | |
| S6* | +Port A_out | | P6 | Ground | |
| S7 | Port A Ground | | P7 | 5 Volts charge | |
| S8 | Port B Ground | | P8* | 5 Volts | |
| S9* | +Port B_in | Diff. input pair | P9* | 5 Volts | |
| S10* | -Port B_in | | P10 | Ground | |
| S11 | Port A Ground | | P11* | Ready LED | Open collector out |
| S12* | -Port B_out | Diff output pair | P12 | Ground | |
| S13* | +Port B_out | | P13 | 12 Volts charge | |
| S14 | Port B Ground | | P14* | 12 Volts | |
| | | | P15* | 12 Volts | |

* - Short pin to support hot plugging

** - Power Disable (T10 Industry Standard) for remote management of the end device.

Allows power cycling / power saving to be controlled by the host via interface pin 3.

Pin P3 of the power connector on the interface may be pulled high, above 1.8V, to shut off power to the drive.

Pulling pin p3 low, below 0.7V (but not negative), allows the drive to run.

NC - No connection in the drive.

11.4.6 SAS transmitters and receivers

A typical SAS differential copper transmitter and receiver pair is shown in Figure 17. The receiver is AC coupling to eliminate ground shift noise.




Figure 17. SAS transmitters and receivers

11.4.7 Power

The drive receives power (+5 volts and +12 volts) through the SAS device connector.

Three +12 volt pins provide power to the drive, 2 short and 1 long. The current return for the +12 volt power supply is through the common ground pins. The supply current and return current must be distributed as evenly as possible among the pins.

Three +5 volt pins provide power to the drive, 2 short and 1 long. The current return for the +5 volt power supply is through the common ground pins. The supply current and return current must be distributed as evenly as possible among the pins.

Current to the drive through the long power pins may be limited by the system to reduce inrush current to the drive during hot plugging.

11.5 Signal characteristics

This section describes the electrical signal characteristics of the drive's input and output signals. See Section 11.4.5 for signal type and signal name information.

11.5.1 Ready LED Out

The Ready LED Out signal is driven by the drive as indicated in the table below.

Table 17 Ready LED Out conditions

| Normal command activity | LED status | |
|--|---|-----|
| Ready LED Meaning bit mode page 19h | 0 | 1 |
| Spun down and no activity | Off | Off |
| Spun down and activity (command executing) | On | On |
| Spun up and no activity | On | Off |
| Spun up and activity (command executing) | Off | On |
| Spinning up or down | Blinks steadily (50% on and 50% off, 0.5 seconds on and off for 0.5 seconds) | |
| Format in progress, each cylinder change | Toggles on/off | |

The Ready LED Out signal is designed to pull down the cathode of an LED. The anode is attached to the proper +3.3 volt supply through an appropriate current limiting resistor. The LED and the current limiting resistor are external to the drive. See the next table for the output characteristics of the LED drive signals.

Table 18 LED drive signal

| State | Test condition | Output voltage |
|---------------|-----------------------------|--------------------------------------|
| LED off, high | 0 V \leq VOH \leq 3.6 V | -100 μ A $<$ IOH $<$ 100 μ A |
| LED on, low | IOL = 15 mA | 0 \leq VOL \leq 0.225 V |

11.5.2 Differential signals

The drive SAS differential signals comply with the intra-enclosure (internal connector) requirements of the SAS standard.

The table below defines the general interface characteristics

Table 19 General interface characteristics

| Characteristic | Units | 3.0Gb/s | 6.0Gb/s | 12.0Gb/s |
|------------------------------------|-------|-----------|-----------|-----------|
| Bit rate (nominal) | Mbaud | 3,000 | 6,000 | 12,000 |
| Unit interval (UI)(nominal) | ps | 333.3 | 166.6 | 83.3 |
| Impedance (nominal, differential) | ohm | 100 | 100 | 100 |
| Transmitter transients, maximum | V | \pm 1.2 | \pm 1.2 | \pm 1.2 |
| Receiver transients, maximum | V | \pm 1.2 | \pm 1.2 | \pm 1.2 |

11.6 SAS-3 Specification compliance

Seagate SAS-3 compatible drives are compliant with the latest SAS-3 Specification (T10/BSR INCITS 519 rev. 06).

The main difference from SAS-2 is the Tx and Rx training that allows the host and drive to adjust the amplitude and emphasis values to the channel. The receiver still employs Decision Feedback Equalizer (DFE) and Feed Forward Equalizer (FFE) circuitry to accomplish this training.

11.7 Additional information

Please contact the Seagate representative for SAS electrical details, if required.

For more information about the Phy, Link, Transport, and Applications layers of the SAS interface, refer to the Seagate *SAS Interface Manual*, part number 100293071.

For more information about the SCSI commands used by Seagate SAS drives, refer to the Seagate *SCSI Commands Reference Manual*, part number 100293068.



Seagate Technology LLC

AMERICAS Seagate Technology LLC 10200 South De Anza Boulevard, Cupertino, California 95014, United States, 408-658-1000

ASIA/PACIFIC Seagate Singapore International Headquarters Pte. Ltd. 7000 Ang Mo Kio Avenue 5, Singapore 569877, 65-6485-3888

EUROPE, MIDDLE EAST AND AFRICA Seagate Technology SAS 16-18 rue du Dôme, 92100 Boulogne-Billancourt, France, 33 1-4186 10 00

Publication Number: 100746003, Rev. B

October 2014